arXiv:1703.11005v2 [cs.DC] 4 Apr 2017

# A simplicial complex model of dynamic epistemic logic for fault-tolerant distributed computing

Eric Goubault* Sergio Rajsbaum†

April 5, 2017

The usual epistemic S5 model for multi-agent systems is a Kripke graph, whose edges are labeled with the agents that do not distinguish between two states. We propose to uncover the higher dimensional information implicit in the Kripke graph, by using as a model its dual, a chromatic simplicial complex. For each state of the Kripke model there is a facet in the complex, with one vertex per agent. If an edge $(u, v)$ is labeled with a set of agents $S$, the facets corresponding to $u$ and $v$ intersect in a simplex consisting of one vertex for each agent of $S$. Then we use dynamic epistemic logic to study how the simplicial complex epistemic model changes after the agents communicate with each other. We show that there are topological invariants preserved from the initial epistemic complex to the epistemic complex after an action model is applied, that depend on how reliable the communication is. In turn these topological properties determine the knowledge that the agents may gain after the communication happens.

## Contents

**1 Introduction** **3**

**2 Distributed computing models** **7**

2.1 Informal overview of fault-tolerant distributed computing theory 7

2.1.1 Overview of distributed systems . . . . . 7

2.1.2 Overview of distributed tasks . . . . . 9

2.1.3 Computability of distributed tasks . . . . . 9

*LIX, Ecole Polytechnique, CNRS, Université Paris-Saclay, 91128 Palaiseau, France goubault@lix.polytechnique.fr

†Instituto de Matematicas, UNAM, Ciudad Universitaria Mexico 04510, Mexico rajsbaum@im.unam.mx

2.2 Formal model . . . . . 10
2.2.1 States, protocols, runs . . . . . 11
2.2.2 Shared memory models . . . . . 13
2.2.3 Models of distributed computation . . . . . 14
2.2.4 Examples of models . . . . . 15

**3 Distributed computability in terms of Kripke frames 17**
3.1 Kripke frames . . . . . 17
3.2 Task solvability . . . . . 19
3.3 Carrier maps as products . . . . . 20

**4 A dynamic epistemic logic semantics for distributed computing 22**
4.1 Dynamic epistemic logic . . . . . 22
4.2 Action models for tasks . . . . . 25
4.3 Action models in distributed computing . . . . . 26

**5 Combinatorial topology and Kripke models 28**
5.1 Simplicial complexes and simplicial models . . . . . 29
5.2 Equivalence between (proper) Kripke models and simplicial models . . . . . 30
5.3 Action models, protocol complexes and task solvability . . . . 33
5.4 Common knowledge and connectivity . . . . . 38

**6 Conclusions 39**

# 1 Introduction

Dynamic epistemic logic (DEL) considers multi-agents systems and studies how their knowledge changes when communication events occur. It extends ordinary epistemic logic by the inclusion of *event models* to describe actions. An epistemic Kripke model represents the knowledge of the agents about an initial situation, and an event model represents their knowledge about the possible event taking place in this situation. The most common situation studied, which gave birth to DEL, is that of a public announcement to all the agents of a formula $\psi$, but there is a general logical language to reason about information and knowledge change [5, 6]. A product update operator defines a Kripke model that results as a consequence of executing actions on the initial Kripke model.

**Overview** An epistemic S5 model is typically used to represent states of a multi-agent system, where edges of the Kripke structure are labeled with the agents that do not distinguish between the two states. In this work we show there are underlying topological invariants induced by the action model, that determine what the agents may know after a communication event takes place.

We argue that while the Kripke model seems to be a one-dimensional structure (a graph), it actually encodes a high dimensional topological object, namely, a *simplicial complex* corresponding, in a precise categorical sense, to the dual of the Kripke structure. Thus, there is a sort of epistemic complex representing the knowledge of the agents about the initial situation, and another epistemic complex, which represents their knowledge obtained through the product with the communication action model.

We show that these complexes indeed carry topological information. In the figure below, $I$ represents an initial epistemic simplicial complex model (equivalent to Kripke model), and $P$ the product with an action model. The action model preserves topological invariants from the initial model $I$ to the complex $P$ after the communication actions have taken place. Actually, how much of the topology is preserved, depends on how reliable is the communication among the agents (see explanation below), and encode the degree at which some knowledge does not evolve.

We explore a class of action models that *fully* preserve the topology of the initial complex, and hence, in a precise sense, yield the *least* information to the agents.

We define another *knowledge goal* action model, that when it is used to make the product with the initial epistemic model, yields an epistemic model $\Delta$ representing what the agents should be able to know, after applying the communication action model. There is sufficient knowledge in $P$ if there exists a (properly defined) morphism $h$ from $P$ to $\Delta$ that makes the diagram commute.

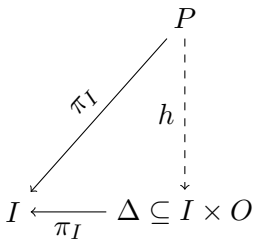


We stress that this is far from telling the whole story. Our formalization of knowledge goals represents only the case of what the processes should learn *once*. Indeed, we provide a precise categorical equivalence between dynamic epistemic Kripke models of tasks and tasks as epistemic complexes (so we have two versions of this diagram, one for Kripke models and one for complex models). A more general situation would study requirements on how knowledge should evolve with time. Also, we concentrate on a very specific setting. The class of communication action models we study, represents all possible ways in which the agents can send to each other messages (but we use shared memory to model the communication pattern), allowing for any pattern of arrivals, including message losses. Roughly speaking, assume a set of agents $A = \{0, \ldots, n\}$ exchange messages once. Each agent sends a single message to every other agent. A communication event defines which messages are delivered and in which order, assuming no partitions happen. Although our setting could be used to study even weaker situations, where messages are lost in a way that there is no communication between two groups of agents. This is, essentially, the *one round* communication action model we consider, and where we can show that the topology is fully preserved, and knowledge gained is minimal (except for of course weaker partitioning situations). Then we study multi-round versions of the action model, and show that also in these the topological invariants are fully preserved.

Finally, let us mention that what we developped here is oriented towards distributed tasks, where problem are specified in terms of input-output (values) relationships. This does not account for inputless problems such as counting problems, timestamps problems etc., for which our framework is going to be generalized in a subsequent paper.

**Main results** As a concrete case study of the theory described above, we consider a multi-agent system representing a distributed computing system, where processes can fail by crashing, and communicate by reading and writing shared variables. Using shared memory instead of message passing

simplifies some of the technical development. Indeed, we have been inspired to consider simplical complexes as the dual of Kripke models from the study of distributed computability of tasks using topology [18]. This approach to analysing distributed computations represents executions as simplicial complexes. As the computation evolves, the topology of the complex changes. A *task* is defined by a relation from an input complex $\mathcal{I}$ specifying the input values to the processes, and an output complex $\mathcal{O}$ specifying values that the processes should decide after communicating with each other. The seminal result in the area is the algebraic topology *set agreement* impossibility, saying that $n+1$ processes cannot agree in at most $n$ different values, *wait-free*, i.e., if the processes are asynchronous and may crash [10, 21, 36]. The categorical equivalence between Kripke models and simplicial complex epistemic models we establish provides, in one direction, a DEL semantics to distributed computability, opening the possibility of reasoning about knowledge change in distributed computing. For instance, the set agreement impossibilty is exposed as the impossibility of gaining sufficient knowledge as required by a decision action model. In the other direction, the connection allows to bring in the topological invariants known in distributed computing, to DEL, and in particular show that knowledge gained after an epistemic action model is intimately related to higher dimensional topological properties.

Section 2 describes the distributed computing model used as a case study. We formalize task solvability for distributed computing in Section 3 in terms of Kripke frames, one of the classical models for multi-modal S5 logics. For this Kripke frame formalization, we note that the classical "carrier map" approach of [18] can be expressed as the existence of a certain commutative diagram involving only morphisms of Kripke frames, see Theorem 1. We then show in Section 4 that these Kripke frames can be used to model a logic of knowledge for agents, and lift Theorem 1 on Kripke models, this is Theorem 3. There is then a simple epistemic interpretation of task solvability : the protocol can only improve knowledge of the agents, and a task is solvable only if it has improved so that to have at least the knowledge formalized by the specification. The other major contribution of this paper is Section 5. We first show that the category of (proper) Kripke frames is equivalent to the category of pure chromatic simplicial complexes, Theorem 4, which is the one in which many results have been developped using combinatorial topology, over the years, see [18]. This also lifts to Kripke models and what we call simplicial models, Theorem 5. From this, we show that the classical protocol complex approach of [18] and our epistemic logic approach are equivalent. The interpretation in terms of knowledge progress and loss of it, has its topological counterpart : there are some topological invariants on the

simplicial models that measure this, extending the intuition that common knowledge is linked to some connectivity conditions. Conclusions and open problems are in Section 6.

Note that we could also have directly defined an interpretation of epistemic logics in (pure chromatic) and later, proved its correctness. We prefered to base our presentation on classical grounds (Kripke models) and show and use an equivalence of categories to end up with a correct interpretation of epistemic logics in simplicial complexes, hence in a geometric framework.

**Related work** Epistemic logic has been used very successfully and many times to analyze and design fault tolerant distributed systems. In this paper we use DEL [7, 14]. Complex epistemic actions can be represented in *action model logic* [6, 14]. Various examples of epistemic actions have been considered, especially *public announcement logic*, a well-studied example of DEL, with many applications in dynamic logics, knowledge representation and other formal methods areas. However, to the best of our knowledge, it has not been used in distributed computing theory, where fault-tolerance is of primal interest. DEL [6, 14] extends epistemic logic through dynamic operators formalizing information change. Plaza [30] first extended epistemic logic to model public announcements, where the same information is transmitted to all agents. Next, a variety of approaches (e.g., [5, 14]) generalized such a logic to include communication that does not necessarily reach all agents. Here, we build upon the approach developed by Baltag et al. [5] employing "action models". We have focused in this paper on the classical semantics of multi-modal S5 logics. The seminal work on the subject, see e.g. [7], has considered topological models. Future work will include the relationship between these topological models and the geometric realization of our simplicial models of Section 5. Another classical model for epistemic logic is the one of Interpreted Systems of Fagin, Halpern, Moses and Vardi (see e.g. [8]). These are related to the semantics we give of our distributed systems, in Section 2.2.1. More importantly, the adjunctions between Kripke frames and Interpreted Systems, developped in [31] are in the same spirit as the equivalence of categories between (proper) Kripke frames and pure chromatic simplicial complexes that we develop in Section 5.2. The structure of pure chromatic simplicial complexes has the advantage to behave well with respect to products with a natural interpretation in terms of action models, and also conveys some important geometrical contents, that we develop in Section 5.2.

Seminal work on knowledge and distributed systems is of course one of the inspirations of the present work, see e.g. [35], as well as the combinatorial

topology approach for fault-tolerant distributed computing, see e.g. [18]. Quite a few results have been achieved using epistemic logic to model fault-tolerant distributed protocols, see e.g. [29] but the authors know no previous work on relating the combinatorial topological methods of [18] with Kripke models. We should note that the formulation of carrier maps as products which is developed in Section 3.3 seems to have been partially observed in [17].

# 2 Distributed computing models

In a *wait-free* model of $n+1$ processes, a process does not use instructions that wait for events in other processes. Wait-free computation has been studied thoroughly e.g. [23, 34], and has turned out to be fundamental in the theory of distributed computing [18]. Simulations and reductions can be used to transfer results to other models e.g. [12, 16, 19]. For instance, a wait-free model can be used as a basis to study models where $t$ processes may crash (and it makes sense for a process to wait until it hears from $n+1-t$ processes). Wait-free models tolerate any number of processes crashing, yet the treatment of failures is substantially simplified, and are mostly taken care implicitly, as we shall see below. Indeed, we consider only models where failures are not detectable in a finite execution (because it is indistinguishable if a process crashed or is just slow); as opposed to synchronous models, where if a message does not arrive by a given time, a process has crashed. We concentrate in models where communication is by read/write shared variables, but the framework can be used to study more powerful communication primitives available in practice e.g. [25], as well as message passing models.

## 2.1 Informal overview of fault-tolerant distributed computing theory

We start by recalling some of the basics of distributing computing, most relevant to this work; several textbooks contain further details [3, 22, 28].

### 2.1.1 Overview of distributed systems

The most basic model we consider is the *one-round read/write* model ($\mathsf{WR}$), e.g. [3]. It consists of $n+1$ processes denoted by the numbers $[n] = \{0, 1, \ldots, n\}$. A process is a deterministic (possibly infinite) state machine. Processes communicate through a shared memory array $\mathsf{mem}[0 \ldots n]$

which consists of $n+1$ single-writer/multi-reader atomic registers. Each process accesses the shared memory by invoking the atomic operations $\mathsf{write}(x)$ or $\mathsf{read}(j)$, $0 \leq j \leq n$. The $\mathsf{write}(x)$ operation is used by process $i$ to write value $x$ to its own register, $i$, and process $i$ can invoke $\mathsf{read}(j)$ to read register $\mathsf{mem}[j]$, for any $0 \leq j \leq n$. In its first operation, process $i$ writes a value to $\mathsf{mem}[i]$, then it reads each of the $n+1$ registers, in an arbitrary order. Such a sequence of operations, consisting of a write followed by all the reads, is abbreviated by $\mathsf{WScan}(x)$.

A *state* of the system consist of the state of each process and the state of the shared memory. An *execution* is defined by an initial state of the system and an interleaving of the operations of the processes; it is an alternating sequence of states and operations. In the most basic case, *every* interleaving of the operations is a possible execution. Namely, we have a fully asynchronous system.

Several models derived from the one-round read/write WR model have been considered in the literature.

- In the *multi-round* version of the WR model, processes can execute any number of write and read operations on the array mem. It is often convenient to assume the program of a process is structured in rounds, each one consisting of a WScan() operation[1]. This is the most commonly considered version, e.g. [3].

- In the *iterated* WR model, processes communicate through a sequence of arrays. They all go through the sequence of arrays $\mathsf{mem}_0$, $\mathsf{mem}_1 \ldots$ in the same order, executing a single WScan() operation on $\mathsf{mem}_r$, for each $r \geq 0$. Namely, process $i$ executes one write to $\mathsf{mem}_r[i]$ and then reads one by one all entries $j$, $\mathsf{mem}_r[j]$, in arbitrary order, before proceeding to do the same on $\mathsf{mem}_{r+1}$. For an overview of iterated models see e.g. [32].

  Notice that the one round WR model is the special case of either the multi-round or the iterated WR models, when the program of all processes consists of a single WScan() operation.

- The *snapshot* versions of the previous, multi-round and iterated models, is obtained by replacing the WScan() by a WSnap() operation, that guarantees that the reads of the $n+1$ registers happen all atomically, at the same time. In this case an execution is an interleaving of write

[1]This assumption is done without loss of generality for computability purposes; to reduce complexity a program would not always be structured this way.

and snap operations by the processes. Snapshot operations can be implemented wait-free and are a very useful abstraction, see e.g. [1].

- The *immediate snapshot* version of the previous multi-round and iterated models is obtained by assuming that the whole WSnap() operation happens "atomically," and hence is called an ImmSnap() operation. More precisely, an execution in such a model is represented by a sequence of concurrency classes, where in each class a set of ImmSnap() operations take place concurrently [2, 10, 36]. Thus, in a concurrency class, the writes by the set of processes participating in the concurrency class occur in arbitrary order, followed by a read to all registers by each of these processes, in arbitrary order. The iterated version has been considered since [11].

### 2.1.2 Overview of distributed tasks

A *decision task* is the distributed equivalent of a function, where each process knows only part of the input, and after communicating with the other processes, each process computes part of the output. For instance, in the *k-set agreement* task [13] each process starts with an input value from some set of input values $V$ with $|V| > k$, and each process has to decide an output value, that is one of the input values, and such that at most $k$ different values are decided; when $k = 1$ we get the *consensus* task. The $k$-set agreement task is solvable (by asynchronous processes communicating using read/write registers) if and only if at most $t$ processes can crash, for $t < k$, see [10, 12, 15, 21, 36].

### 2.1.3 Computability of distributed tasks

A central concern in distributed computability is studying which tasks are solvable in a given distributed computing model (as determined by e.g. the type of communication mechanism available and the reliability of the processes). Early on it was shown that consensus is not solvable even if only one process can fail by crashing, when the processes are asynchronous and they communicate by message passing [15], or even by writing and reading a shared memory [27]. A graph theoretic characterization of the tasks solvable in the presence of at most one process failure appeared soon after [9].

The *asynchronous computability theorem* [21] exposed that moving from tolerating one process failure, to any number of process failures, yields a

characterization of the class of decision tasks that can be solved in a wait-free[2] manner by asynchronous processes based on simplicial complexes, which are higher dimensional versions of graphs. In particular, $n$-set agreement is not wait-free solvable, with $n+1$ processes [10, 21, 36].

Computability theory through combinatorial topology has evolved to encompass non-independent process failures, arbitrary malicious failures, synchronous and partially synchronous processes, and various communication mechanisms [18]. Still, the original wait-free model of the asynchronous computability theorem, where crash-prone processes that communicate wait-free by writing and reading a shared memory is fundamental. Topological techniques are derived in this model, and then extended to other models, e.g. [20]. Also, the question of solvability in other models (e.g. $t$ crash failures, for $1 \leq t \leq n$), can in many cases be reduced to the question of wait-free solvability ($t = n$), as shown in [12] and [19].

More specifically, in the *AS model* of [18] each process can write its own location of the shared-memory, and it is able to read the entire shared memory in one atomic step, called a *snapshot*. The characterization is based on the *protocol complex,* which is a geometric representation of the various possible executions of a protocol. Simpler variations of this model have been considered. In the *immediate snapshot* (IS) version [2, 10, 36], processes can execute a combined write-snapshot operation. The *iterated immediate snapshot* (IIS) model [11] is even simpler to analyze, and can be extended (IRIS) to analyze partially synchronous models [33]. Processes communicate by accessing a sequence of shared arrays, through immediate snapshot operations, one such operation in each array. The success of the entire approach hinges on the fact that the topology of the protocol complex of a model determines critical information about the solvability of the task and, if solvable, about the complexity of solution [24].

All these snapshot models, AS, IS, IIS and IRIS can solve exactly the same set of tasks. However, the protocol complexes that arise from the different models are structurally distinct. The combinatorial topology properties of these complexes have been studied in detail, providing insights for why some tasks are solvable and others are not in a given model.

## 2.2 Formal model

We adapt the model of [29] (in turn following the style of [35]), for the wait-free case.

[2]When any number of processes may crash, the algorithm run by a process must be *wait-free,* because there is no reason for it to wait for an event occurring in another process.

### 2.2.1 States, protocols, runs

There is a fixed finite set of processes, $0, 1, 2, \ldots, n$, and an *environment,* $e$, which is used to model aspects of the system that are not modeled as being part of the activity or state of the processes. We will model the shared memory as being part of the environment's state. In addition, we will assume that various nondeterministic choices such as various delays are actions performed by the environment. For every $i \in \{e, 0, 1, \ldots, n\}$, we assume there is a set $L_i$ consisting of all possible *local states* for $i$. The set of *global states,* which we will simply call *states*, will consist of $\mathcal{G} = L_e \times L_0 \times \cdots \times L_n$.. We denote by $x_i$ the local state of $i$ in the state $x$.

In a given setting, every $i \in \{e, 0, 1, \ldots, n\}$ is associated with a nonempty set $\mathsf{ACT}_i$ of possible actions. Intuitively, these model shared-memory operations or local computations the process may perform. The environment is in charge of *scheduling* the processes. A *scheduling action* is a set $\mathsf{Sched} \subseteq \{0, \ldots, n\}$ of the processes that are scheduled to move next. We assume the existence of a set $\mathsf{act_e}$ describing the aspects of the environment's actions that handle everything other than the scheduling of processes. Without loss of generality we will assume that an environment's action (an element in $\mathsf{ACT}_e$) is a pair $(\mathsf{Sched}, a)$, where $\mathsf{Sched}$ is a scheduling action and $a \in \mathsf{act_e}$. A *joint action* is a pair $\bar{\mathsf{a}} = (sa, \mathbf{a})$, where $sa = (\mathsf{Sched}, a)$ is in $\mathsf{ACT}_e$, and $\mathbf{a}$ is a function with domain $\mathsf{Sched}$ such that $\mathbf{a}(i) \in \mathsf{ACT}_i$ for each $i \in \mathsf{Sched}$. Thus, $\bar{\mathsf{a}}$ specifies an action for the environment (via $\mathbf{a}$) and an action for every process that is scheduled to move. We define the set of joint actions by $\overline{\mathsf{ACT}}$. Clearly, $\overline{\mathsf{ACT}}$ is determined by a collection of action sets $\{\mathsf{ACT}_i\}_{i=1,\ldots,n}$ and a set $\mathsf{act_e}$. Joint actions are the events that cause the global state to change into a new state. This is formally captured by the notion of a *transition function*, which is a function $\tau : \mathcal{G} \times \overline{\mathsf{ACT}} \to \mathcal{G}$ from global states and joint actions to global states, describing how a joint action transforms the global state.

Processes follow a *deterministic protocol*, a function $P_i : L_i \to \mathsf{ACT}_i$ specifying the action that $i$ is ready to perform in every state of $L_i$. The environment however follows a *nondeterministic protocol*, a function $P_e : L_i \to 2^{\mathsf{ACT}_i} \setminus \emptyset$ specifying for every state of $L_e$ a nonempty set of actions, one of which $e$ must perform in that state.

A run is defined by a sequence of global states and the joint actions that cause the transitions among them. Notice that once we fix a deterministic protocol $D = (D_0, \ldots, D_n)$ for the processes, an action $sa = (\mathsf{Sched}, a)$ of the environment uniquely determines the joint action $\bar{\mathsf{a}} = (sa, \mathbf{a})$ that will be performed in a given (global) state $x$: the set $\mathsf{Sched}$ determines

the processes that participate in the joint action, and $\mathbf{a}(i) = D_i(x_i)$ for each $i \in \mathsf{Sched}$. Formally, we model a *run* over $\mathcal{G}$ and $\mathsf{ACT}_e$ as a pair $R = (r, \alpha)$, where $r : N \to \mathcal{G}$ is a function from the natural numbers to $\mathcal{G}$ defining sequence of states of $\mathcal{G}$, and $\alpha : N \to \mathsf{ACT}_e$ defines a corresponding sequence of environment actions. The intuition will be that the joint action caused by $\alpha(k)$ and the underlying protocol leads us from a state $r(k)$ to a state $r(k+1)$. As we will see later on, once we fix a model of computation and a protocol $D$ for the processes to follow, there will be additional conditions relating the sequences $r$ and $\alpha$. These conditions guarantee, for example, that the actions recorded by $\alpha$ do indeed cause the transitions among the corresponding states recorded by $r$. The state $r(0)$ is called the *initial state* of the run $R$. We denote by $r(k)_i$ (resp., $r(k)_e$) the local state of process $i$ (resp., of the environment) in $r(k)$.

An *execution* is a subinterval of a run, starting and ending in a state, as described next. For a run $R = (r, \alpha)$ and a pair $m \leq m'$ where $m$ is finite and $m'$ is finite or infinite, we denote by $R[m, m']$ the execution starting at the state $r(m)$ and ending in $r(m')$ and behaving as $R$ does between them. Formally, $R[m, m'] = (\sigma, \beta)$ where $\sigma$ has domain $[0, m' - m]$ and $\beta$ has domain $[0, m' - m - 1]$, and they satisfy $\sigma(k) = r(m+k)$ and $\beta(k) = \alpha(m+k)$ for every $k$ in their respective domains. Notice that, in principle, the same execution can occur in different runs, and for that matter even at different times. A *suffix* of a run $R$ is an execution of the form $R[m, \infty]$ for some finite $m$; similarly, a *prefix* of $R$ is an execution of the form $R[0, m]$.

Given an execution $R$ (possibly consisting of just one state), let us denote by $R \odot sa$ the execution that results from extending $R$ by having the environment perform the action $sa$. In models in which performing a joint action at a state results in a unique next state (which will invariably be the case in this paper), every run of a deterministic protocol $D$ can be represented in the form $x \odot sa_1 \odot sa_2 \odot \cdots$ where $x$ is an initial state and $sa_i$ is an environment action, for every integer $i \geq 0$.

A *system* $\mathcal{R}$ is a set of runs. With respect to a system $\mathcal{S}$, a state $y$ is said to *extend* the state $x$ if there is a finite execution in some run of $\mathcal{S}$ that starts in $x$ and ends in $y$. A run $R$ is said to *contain* a state $x$ if $x$ is one of the global states in $R$. For conciseness, we will use terminology such as *a state $x$ of $\mathcal{S}$*, when we mean a state $x$ appearing in a run of $\mathcal{S}$, or *an initial state of $\mathcal{S}$*, when we mean a state appearing as an initial state in a run of $\mathcal{S}$. By convention, $x$ extends $x$ for every state $x$ of $\mathcal{S}$.

### 2.2.2 Shared memory models

The standard asynchronous shared-memory model is well known, and detailed formal descriptions can be found in textbooks such as [3, 22]. We now briefly review the features of the model that are relevant for the analysis presented in this section.

We assume the standard asynchronous shared-memory model where $n+1$ processes, $n \geq 2$, communicate by reading and writing to single-writer/multi-reader, shared variables, and any number of processes can crash. A (global) *state* $x$ of the system is a tuple specifying a local state $x_i$ for every process $i$, and the state of the environment, which in this case consists of the assignment of values to the shared variables, as well as the set of pending shared-memory operations, and the set of pending reports for read operations that have been recorded (the value has been read) but not yet reported to the reading process. The pending operations are the read and write operations that have been issued for these variables and have not yet taken effect.

The sets $\mathsf{ACT}_i$ and $\mathsf{ACT}_e$ of the actions of the processes and the environment are defined as follows. A process performs an action only when it is scheduled to move. This action is either a read of a shared variable (belonging to it or to some other process), or a write to one of its own variables. For simplicity we do not include local operations, we assume they are part of the read and write operations. An action of the environment can have one of three forms: (a) scheduling a process to move—resulting in the process performing an action, (b) performing a pending shared-memory operation, or (c) reporting the value read in a recorded read operation to the reading process.

A run of a given deterministic protocol $D$ in this model is a run $R = (r, \alpha)$ satisfying the following:

(i) $r(0)$ belongs to a set of possible initial states,

(ii) each process follows its protocol, and every pair of consecutive states are related according to the operations that take place as scheduled by the environment.

If in addition

(iii) the read and write actions issued are appropriately are eventually serviced by the environment,

then we say that the run is *admissible*.

Let $\mathcal{S}(D)$ be the system consisting of the set of all admissible runs of $D$, where each process is scheduled to move exactly the same number of times,

$N$. If a task is solvable, it is solvable in a finite number of operations, $N$. Alternatively, when a task is not solvable, there is no $N$ for which it is solvable. We may also consider systems that are strictly contained in $\mathcal{S}(D)$, where although all processes execute $N$ operations, not all interleavings are included.

Recall that $R \odot sa$ is the execution that results from extending $R$ by having the environment perform the action $sa$. In a shared memory model performing a joint action at a state results in a unique next state and assuming the environment serves immediately every read and write operation, every run of a deterministic protocol $D$ can be represented in the form $x \odot sc_1 \odot sc_2 \odot \cdots$ where $x$ is an initial state and $sc_i$ is a scheduling action, for every integer $i \geq 0$. That is, it is sufficient to state which is the set of processes that are scheduled to move in a state, to determine the action of the environment, the action of the processes, and what will be the next state.

We continue by defining models of distributed systems in a more general manner, that can capture snapshot and immediate snapshot models.

### 2.2.3 Models of distributed computation

We define a generic model of computation, that can be used in situations other than shared memory. A *model of distributed computation* is determined by sets $L_i$, $i \in \{e, 0, 1, \ldots, n\}$, of local states for the processes and the environment, and corresponding sets of actions $\mathsf{ACT}_i$, for every $i \in \{e, 1, 2, \ldots, n\}$, and by a tuple $M = (\mathcal{G}_0, P_e, \tau, \Psi)$, where the following hold:

- $\mathcal{G}_0 \subseteq \mathcal{G}$ is called the set of *initial states*. The identity of $\mathcal{G}_0$ will depend on the type of analysis for which the model is introduced. When we focus on a particular problem such as consensus, $\mathcal{G}_0$ is the set $\mathsf{Con}_0$ of initial states for consensus.

- $P_e$ is a (nondeterministic) protocol for the environment.

- $\tau$ is a transition function.

- $\Psi$ is a set of runs over $\mathcal{G}$ and $\mathsf{ACT}_e$, such that for every pair of runs $R$ and $R'$ that have a suffix in common, $R \in \Psi$ if and only if $R' \in \Psi$. The set $\Psi$ is called the set of *admissible* runs in the model. This is a tool for specifying fairness properties of the model. For example, properties such as "every message sent is eventually delivered" or "every process moves infinitely often" are enforced by allowing as admissible only runs in which these properties hold. The condition we have on $\Psi$ being

determined by the suffixes of its runs ensures that admissibility depends only on the infinitary behavior of the run.

We say that a run $R = (r, \alpha)$ is a *run of the protocol* $D = (D_1, \ldots, D_n)$ *in* $M = (\mathcal{G}_0, P_e, \tau, \Psi)$ when

(i) $r(0) \in \mathcal{G}_0$, so that $R$ begins in a legal initial state according to $M$,

(ii) $\alpha(k) \in P_e(r(k)_e)$ for all $k$,

(iii) $r(k+1) = \tau(r(k), (\alpha(k), \mathbf{a}^k))$ holds for all $k \geq 0$, where the domain of $\mathbf{a}^k$ is the set $\mathsf{Sched}$ in $\alpha(k)$, and $\mathbf{a}^k_i = D_i(r(k)_i)$ for every $i \in \mathsf{Sched}$,[3] and

(iv) $R \in \Psi$, so that $R$ is admissible.

Condition (ii) implies that the environment's action at every state of $R$ is legal according to its protocol $P_e$, and condition (iii) states that the state transitions in $R$ are according to the transition function $\tau$, assuming that the joint action is the one determined by the environment's action and the actions that the protocol $D$ specifies for the processes that are scheduled to move. A run satisfying properties (i)–(iii) but not necessarily the admissibility condition (iv) is called *a run of* $D$ *consistent with* $M$. It is a run in which the initial state and local transitions are according to $D$ and $M$, but the admissibility conditions imposed by $\Psi$ are not necessarily satisfied.

The notions of models and protocols give us a way of focusing on a special class of systems, resulting from the execution of a given protocol in a particular model. We denote by $\mathcal{S}_N(D, M, I)$ the system $\mathcal{R}$, where $\mathcal{R}$ is the set of all executions of protocol $D$ in the model $M$ that start in initial states from a set $I$, where $I \subseteq \mathcal{G}_0$, and each process executes exactly $N$ operations. When (any of) $D, M$ or $N$ are clear from the context, we may simply write $\mathcal{S}(I)$.

### 2.2.4 Examples of models

Consider a set of initial states $\mathcal{G}_0$, where the initial state of a process $i$ is a pair $(i, v)$, and $v \in \{0, 1\}$ represents an input value. The environment is always in the same initial state, $\epsilon$.

[3] Given this choice, any deviations of a process from the protocol, as may happen in a model with malicious failures, will need to be modeled as resulting from the environment's actions. The behavior of faulty processes in such a case will be controlled by the environment.

**IS and IIS models** Consider the iterated immediate snapshot IIS model (Section 2.1.1) where processes communicate through a sequence of arrays $\mathsf{mem}_0$, $\mathsf{mem}_1 \ldots$. The protocol $D$ that the processes follow consists of executing a single WScan() operation on $\mathsf{mem}_r$, for each $r \geq 0$. Namely, process $i$ executes one write to $\mathsf{mem}_r[i]$ and then reads one by one all entries $j$, $\mathsf{mem}_r[j]$, in arbitrary order, before proceeding to do the same on $\mathsf{mem}_{r+1}$.

The set of all $N$-step executions $\mathcal{R}$ of protocol $D$ in the IIS model that start in initial states $\mathcal{G}_0$ can be obtained by applying schedules of the following form. Every execution in $\mathcal{R}$ is of the form $x \odot sc_1 \odot sc_2 \odot \cdots \odot sc_R$ where $x \in \mathcal{G}_0$ and $sc_i$ is a *block scheduling action* for every integer $1 \leq i \leq R$. A block action is an ordered partition of the set $A = \{0, \ldots, n\}$, consisting of *concurrency classes* $[s_0, \ldots, s_k]$, the $s_i$ are non-empty, disjoints subsets of $A$, whose union is $A$. Thus, $0 \leq k \leq n$. When $k = n$ the execution is fully sequential (processes take immediate snapshots one after the other), and when $k = 0$ it is fully concurrent (they all execute an immediate snapshot concurrently).

To apply the block action $[s_0, \ldots, s_k]$ to initial state $x$, and obtain state $x \odot [s_0, \ldots, s_k]$, the environment schedules the processes in the following order, to execute their read and write operations on $\mathsf{mem}_0$. It first schedules the processes in $s_0$ to execute their write operations, and then it schedules them to execute their read operations (the specific order among writes is immaterial, and so is the case for the reads). Then the environment repeats the same for the processes in $s_1$, scheduling first the writes and then the reads, and so on, for each subsequent concurrency class $s_i$. Notice that the read and write operations of block action $sc_i$ are applied to memory $\mathsf{mem}_i$.

Consider the composition of the block actions, $sc_1 \odot sc_2 \odot \cdots \odot sc_R$. Let $IIS_R$ denote the set of all such composition of block actions. That is,

$$sc_1 \odot sc_2 \odot \cdots \odot sc_R \in IIS_R$$

if and only if each $sc_i$ is an ordered partition of $A$, of the form $[s_0, \ldots, s_k]$. Then, any execution of $\mathcal{R}$ of protocol $D$ in the IIS model can be obtained by applying a composed block action of $IIS_R$ to an initial state in $\mathcal{G}_0$. In other words, $\mathcal{R}$ can be seen as the product of $\mathcal{G}_0$ and $IIS_R$.

The non-iterated immediate snapshot IS model (Section 2.1.1) is defined in a similar way, except that processes communicate through a single array mem. The protocol $D$ that the processes follow consists of executing a WScan() operations repeatedly on mem. Namely, process $i$ executes one write to $\mathsf{mem}[i]$ and then reads one by one all entries $j$, $\mathsf{mem}[j]$, in arbitrary order.

Every execution in $\mathcal{R}$ is of the form $x \odot sc_1 \odot sc_2 \odot \cdots \odot sc_R$ where $x \in \mathcal{G}_0$, but now each $sc_i$ is a concurrency class of processes, to be scheduled as before: the environment schedules the processes in the following order, to execute their read and write operations on mem. It first schedules the processes in $sc_1$ to execute their write operations, and then it schedules them to execute their read operations, on the initial state, to obtain a new state, then repeats the same schedule with processes on $sc_2$, and so on.

The only condition is that at the end of the execution all processes have been scheduled the same number of times, $N$, when applying the composition of the block actions, $sc_1 \odot sc_2 \odot \cdots \odot sc_R$. Then, $IS_R$ denotes the set of all such composition of block actions. That is,

$$sc_1 \odot sc_2 \odot \cdots \odot sc_R \in IS_R$$

if and only if each $sc_i$ is a concurrency class of $A$, and each process appears $N$ times over all concurrency classes in the execution.

An interesting example of a schedule by the environment yields a *solo execution* for process $i$, in both the IS and IIS models. In the IS model, it schedule all $N$ first concurrency classes with a singleton set containing $i$, and then schedules all other processes in any order. In the IIS model, it schedules $i$ first, alone, in every one of the $N$ block actions.

# 3 Distributed computability in terms of Kripke frames

The distributed computing modeling in Section 2 is based on executions. Here we move to the orthogonal perspective, of considering sets of states.

## 3.1 Kripke frames

We define three Kripke frames, one for the input states, one for the protocol states after $N$ steps, and one for the output states. Also, two types of relations between Kripke frames: morphism (that will later on correspond to simplicial map) and carrier morphism (that will correspond to carrier map).

In this paper a state of a Kripke frame consists of the local states of the processes (sometimes including the state of the environment, but usually not), and agents correspond to processes in a distributed model. Whenever the states of a Kripke frame have this meaning, we use the following accessibility relation. Two (global) states $u, v \in S$ are defined to be *indistinguishable* by

$a$, $u \sim_a v$, if and only if the state of process $a$ is the same in $u$ and in $v$. Notice that $u \sim_a v$ defined this way is indeed an equivalence relation.

**Definition 1** (Kripke frame)**.** *Assume a set $A = \{a_0, a_1, \ldots a_n\}$ of $n+1$ agents. A* Kripke frame *$M = \langle S, \sim^A \rangle$ consists of a finite set of* states*, $S$, and a function $\sim^A$, yielding for every $a \in A$ an equivalence* $a$-accessibility relation $\sim_a \subseteq S \times S$.

Intuitively, two states in a Kripke structure are connected with an edge labeled with agent $a$ if such states are "indistinguishable" by $a$. We sometimes view $M$ as a graph whose edges are labeled with the agents that do not distinguish between the two states of the edge (implicitly, every vertex has a self-loop labeled with all agents). We will be interested in *proper* Kripke frames (also known as frames satisfying the identity intersection property in [31]), where any two states are distinguishable by at least one agent.

**Definition 2** (Morphism of Kripke frames)**.** *Let $M = \langle S, \sim^A \rangle$ and $N = \langle T, \sim^A \rangle$ be two Kripke frames. A* morphism *of Kripke frame $M$ to $N$ is a function $f$ from $S$ to $T$ such that for all $u, v \in S$, for all $a \in A$, $u \sim_a v$ implies $f(u) \sim_a f(v)$.*

These morphisms are known also as weak morphisms, see e.g. [31] as a more classical notion of morphism is generally the one of $p$-morphism (or bounded morphism, see e.g. [31] again), which we will not be using in the paper.

We call $\mathcal{K}$ the category of Kripke frames, with morphisms of Kripke frames as defined above. Note that this category enjoys many interesting properties, among which the fact that cartesian products exist (Lemma 1) :

**Definition 3.** *Let $M = \langle S, \sim^A \rangle$ and $N = \langle T, \sim^A \rangle$ be two Kripke frames, and define $\langle U, \sim^A \rangle$ as follows : states $U$ are pairs $u = (s, t)$ of states $s \in S$ and $t \in T$ and the accessibility relation is defined as $(s, t) \sim_a (s', t')$ if and only if $s \sim_a s'$ and $t \sim_a t'$. We call $U$ the product Kripke frame.*

Now this product is indeed the categorical product :

**Lemma 1.** *The Kripke frame defined in Definition 3 is the cartesian product, denoted by $M \times N$, in the categorical sense, of $M$ with $N$, coming with projections $\pi_M : M \times N \to M$ and $\pi_N : M \times N \to N$, which are morphisms of Kripke frames.*

Let $M = \langle S, \sim^A \rangle$ and $N = \langle T, \sim^A \rangle$ be two Kripke frames. A *carrier morphism* of Kripke frame $M$ to $N$ is a relation $\Phi$ from $S$ to $T$ such that for

every $s \in S$, $\Phi(s) \subseteq T$, and for all $u, v \in S$, there exist $u' \in f(u), v' \in f(v)$, with $u' \sim_a v'$, for each $a \in A$ for which $u \sim_a v$. Notice that $\Phi(s)$ is non-empty, for each $s \in S$.

It easy is to check that both morphism and carrier morphism compose, in particular we use the following, where $\mathcal{I} = \langle \mathcal{G}_0, \sim^A \rangle$, $\mathcal{P} = \langle \mathcal{G}_P, \sim^A \rangle$ and $\mathcal{O} = \langle \mathcal{G}_f, \sim^A \rangle$ are arbitrary Kripke frames.

Let $M = \langle S, \sim^A \rangle$ and $N = \langle T, \sim^A \rangle$ be two Kripke frames, $f$ a morphism, and $\Phi$ a carrier morphism from $S$ to $T$. We say $f$ *respects* $\Phi$ if $f(s) \in \Phi(s)$, for every $s \in S$. Similarly, $\Phi$ *respects* a carrier morphism $\Delta$ from $S$ to $T$ if $\Phi(s) \subseteq \Delta(s)$, for every $s \in S$. When $f$ respects $\Phi$ we may also say that $f$ *is carried* by $\Phi$.

### 3.2 Task solvability

In the following definition of a task, the input states $\mathcal{G}_0$, are such that in every initial state, the environment is in the same state, say $\epsilon$. Thus the inputs are encoded in the initial states of the processes only (and the shared memory does not contain inputs to be used by the processes), and we may as well disregard the environment state from any initial state. Similarly, there is a set of output states, $\mathcal{G}_f$, with no environment state (or equal to $\epsilon$), that represent what each process should produce as output value in an execution. The relation $\sim^A$ consists of the indistinguishability relations defined above.

**Definition 4** (Task)**.** *A* task *is a triple $\mathcal{T} = \langle \mathcal{G}_0, \mathcal{G}_f, \Delta \rangle$. The set of initial initial states $\mathcal{G}_0$, is such that in every initial state, the environment is in the same state, and $\mathcal{G}_f$ is a set of states, with the same environment state. Then, $\Delta$ is a carrier morphism of Kripke frame $I = \langle \mathcal{G}_0, \sim^A \rangle$ to Kripke frame $O = \langle \mathcal{G}_f, \sim^A \rangle$.*

Note that we can view a task as in the definition above as a sub-Kripke frame $Z$ of the product frame $\mathcal{G}_0 \times \mathcal{G}_f$, by just taking the states of $\mathcal{G}_0 \times \mathcal{G}_f$ that are related by carrier map $\Delta$, within $\mathcal{G}_0 \times \mathcal{G}_f$, with the induced accessibility relations. In fact, $\Delta$, as a carrier map, of some state $s \in \mathcal{G}_0$ is just the image by the second projection $\pi_{\mathcal{G}_f}$ of the sub-frame of $\mathcal{G}_0$ generated by states $g$ such that $\pi_{\mathcal{G}_0}(g) = s$. In the sequel we identify $\Delta$ with $Z$, for simplicity's sake.

Consider a model of distributed computation $M = (\mathcal{G}_0, P_e, \tau, \Psi)$, with initial states $\mathcal{G}_0 \subseteq \mathcal{G}$, a deterministic protocol $D$, and fix an integer $N$. Recall that $\mathcal{S}_N(D, M, \mathcal{G}_0)$ contains the $N$-admissible executions, where each process executes the same number of actions.

The *N-protocol Kripke frame* $P = \langle S, \sim^A \rangle$ consists of states $S$, where $s \in S$, is obtained by removing the environment state from the last state in every execution in this system $\mathcal{S}_N(D, M, \mathcal{G}_0)$.

**Lemma 2.** *Let $I = \langle \mathcal{G}_0, \sim^A \rangle$ be the input Kripke frame and $P = \langle S, \sim^A \rangle$ be the $N$ step protocol Kripke frame. Define $\mathcal{P}$ to be the relation from $\mathcal{G}_0$ to $S$ sending each state $s \in \mathcal{G}_0$ to the states at the end of executions with initial state $s$. Then $\mathcal{P}$ is a carrier morphism from $I$ to $P$.*

**Definition 5** (Solving a Task)**.** *The protocol $D$ with carrier morphism $\mathcal{P}$ solves task $\mathcal{T} = \langle \mathcal{G}_0, \mathcal{G}_f, \Delta \rangle$ in $N$ steps if there exists a morphism, $f$ from $P = \langle \mathcal{G}_P, \sim^A \rangle$ to $O = \langle \mathcal{G}_f, \sim^A \rangle$ such that the composition of $\mathcal{P}$ and $f$ is a carrier morphism respecting $\Delta$.*

Let us discuss this definition. First, assume the protocol $D$ with carrier morphism $\mathcal{P}$ solves task $\mathcal{T} = \langle \mathcal{G}_0, \mathcal{G}_f, \Delta \rangle$ in $N$ steps with a morphism, $f$ from $P = \langle Gp, \sim^A \rangle$ to $O = \langle Gf, \sim^A \rangle$. If $f$ send a state $s$ to an output state $f(s)$ then indeed each process $a$ can decide (operationally, in its program) the value that corresponds to its local state in $f(s)$: (i) the value is a function only of its local state, and (ii) it is consistent in the sense that if $a$ has the same local state in two states $u, v$ then the decision of $a$ are required to be the same by $f$.

Second, these decisions are respecting the task specification, because if we consider an initial state $s_0$, then any execution starting in $s_0$ ends in a state in $s \in \mathcal{P}(s_0)$, which is then mapped to a state $t$ by $f(s)$, with $t \in \Delta(s_0)$.

Finally, if such a carrier morphism $f$ does not exist, then it is impossible to solve the task in $N$ steps by a deterministic protocol $D$ in model $M$, because any such protocol would actually be defining a carrier morphism, because the decision of a protocol are based only on the local states of the processes after $N$ steps.

### 3.3 Carrier maps as products

Fix a model of distributed computation $M = (\mathcal{G}_0, P_e, \tau, \Psi)$, with initial states $\mathcal{G}_0 \subseteq \mathcal{G}$, and a deterministic protocol $D$, with some integer $N$. The set $\mathcal{S}_N(D, M, \mathcal{G}_0)$ of $N$-admissible executions, where each process executes the same number of actions can be seen as a product of $\mathcal{G}_0$ and a set of composed scheduling actions $R$ where each composed block action is of the form $sc_1 \odot sc_2 \odot \cdots \odot sc_R$. If $IIS_R$ denotes the set composition of block actions, $sc_1 \odot sc_2 \odot \cdots \odot sc_R \in IIS_R$ we may define the product $S = \mathcal{G}_0 \times IIS_R$ to be equal to the set of states at the end of all executions starting in a state

in $\mathcal{G}_0$ after applying the actions of an element in $IIS_R$. Given that this is a set $S$ of (global) states, we have the induced indistinguishability relation $\sim^A$, and a Kripke frame, $P = \langle Gp, \sim^A \rangle$. This is of course the protocol Kripke frame which are the states at the end of all $N$-admissible executions starting in $\mathcal{G}_0$.

More generally, for model $M = (\mathcal{G}_0, P_e, \tau, \Psi)$, protocol $D$ and $N$, we can consider the set of all possible composed scheduling actions of the environment $AC$, each one determined by the order in which the environment schedules the processes. Then the protocol Kripke frame can be written as $P = \langle S, \sim^A \rangle$, where $S = \mathcal{G}_0 \times AC$ which is indeed the cartesian product defined in Lemma 1 between frame $\mathcal{G}_0$ and the action frame $AC$. Because the protocol is defined as a product, we have the following simplification for task solvability :

**Theorem 1.** *In the category of Kripke frames, task solvability in the sense of Definition 5 is equivalent to the existence of a morphism $h : P \to \Delta$ such that $\pi_I \circ h$ is equal to $\pi_I : P \to I$, i.e. the following diagram commutes :*

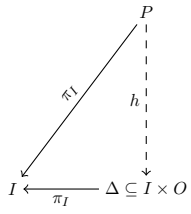


*Proof.* Suppose we have task solvability in the sense of Definition 5. Hence we have a morphism $f$ and carrier maps (identified with thick arrows below) $\mathcal{P}$ such that the following diagram commutes "up to inclusion" : F

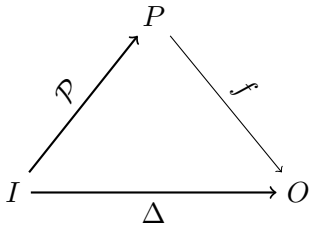


First, we note, as observed after Definition 4 that the carrier map below is a composite of the inverse image of $\pi_{\mathcal{G}_0}$ with $\pi_{\mathcal{G}_f}$. Similarly, as $P$ is a product frame, the carrier map $\mathcal{P}$ is induced by the structural map $\pi_{\mathcal{G}_0} : P \to \mathcal{G}_0$. Hence task solvability is equivalent to the commutation of the diagram below, where $p$ is any section of the projection $\pi_{\mathcal{G}_0} : \Delta \to \mathcal{G}_0$, and $q$ is any section of the projection $\pi_{\mathcal{G}_0} : P \to \mathcal{G}_0$ :

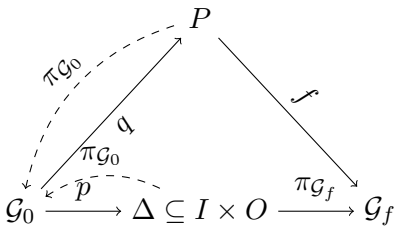


Note that by the universal property of cartesian products, this is equivalent to the existence of a weak morphism $h$ from $P(I)$ to $I \times O$, which furthermore has to be in values in $\Delta \subseteq I \times O$, such that the following diagram commutes :

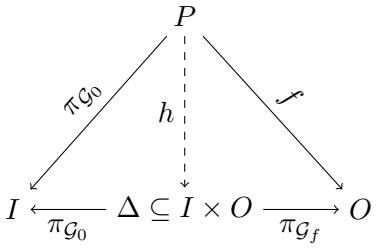


Hence the result. □

# 4 A dynamic epistemic logic semantics for distributed computing

We consider a Kripke model that represents the possible initial states in the model, and use dynamic epistemic logic, to represent the environment as an action model transforming the initial model into Kripke model representing the epistemic state at the end of $N$ steps of protocol $D$. Furthermore, we express desired epistemic states to be reached, by specifying a task also as an action model.

## 4.1 Dynamic epistemic logic

We adhere to the notation of [14]. Let $AP$ be a countable set of *atomic propositions* (i.e., *propositional variables*). The set of *literals* over $AP$ is $\mathrm{Lit}(AP) = AP \cup \{\neg p \mid p \in AP\}$. The *complement* of a literal $p$ is defined by $\overline{p} = \neg p$ and $\overline{\neg p} = p$, $\forall p \in AP$. If $X \subseteq \mathrm{Lit}(AP)$, then $\overline{X} = \{\overline{\ell} \mid \ell \in X\}$; $X$ is *consistent* iff $\forall\ \ell \in X$, $\overline{\ell} \notin X$; and $X$ is *$AP$-maximal* iff $\forall\ p \in AP$, either $p \in X$ or $\neg p \in X$.

**Example 1.** *Consider distributed protocols that manipulate boolean values, i.e. where for all $i \in \{0, \ldots, n\}$, the local values that agents $a_i$ can take are $L_i = \{0, 1\}$. AP is the set $\{l_i \mid i = 0, \ldots, n\}$ of predicates whose informal interpretation is $l_i$ is true when agent $a_i$ holds local value 1 (otherwise, it holds 0).*

**Definition 6** (syntax)**.** *Let AP be a countable set of propositional variables and A a set of agents. The language $\mathcal{L}_K$ is generated by the following BNF grammar:*

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi$$

**Definition 7** (Semantics of formulas)**.** *Epistemic state $(M, s)$ with $M = \langle S, \sim, V\rangle$, and $\varphi \in \mathcal{L}_K(A, AP)$. The satisfaction relation, determining when a formula is true in an epistemic state, is defined as:*

$M, s \models p$ *iff* $p \in L(s)$
$M, s \models \neg\varphi$ *iff* $M, s \not\models \varphi$
$M, s \models \varphi \wedge \psi$ *iff* $M, s \models \varphi$ *and* $M, s \models \psi$
$M, s \models K_a\varphi$ *iff* *for all* $s' \in S : s \sim_a s'$ *implies* $M, s' \models \varphi$

shorthand:

$$E_B\varphi = \bigwedge_{b \in B} K_b\varphi$$

Hence, an agent $a$ is said to know an assertion in a state $(M, s)$ iff that assertion of true in all the states it considers possible, given $s$. Therefore, the *knowledge* $K_a$ of an agent $a$ with respect to a state $s$ is the set of formulas which are true in all states $a$-accessible from $s$.

**Example 2.** *Consider boolean tasks where both input and output values to the processes are taken from the set $\{0, 1\}$ as in Example 1 (and we use the same notations as in this example). Consider now the following Kripke model, where the squiggly arrows picture the map L, associating to each state of the underlying Kripke frame, AP-maximal sets of literals :*

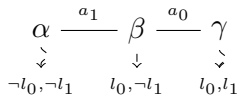


*We have now the following knowledge, known to agents :*

- *In $\alpha$ : $K_0\neg l_1$ and $K_0\neg l_0$ ($a_0$ knows both its value, which is 1, and the value of $a_1$, which is 1) ; $\neg K_1\neg l_0$ and $K_1\neg l_1$ ($a_1$ does not know if the value of $a_0$ is 1, but knows its own value, which is 1)*

- *In $\beta$ : $\neg K_0 \neg l_1$ and $K_0 l_0$ ($a_0$ knows its value, which is 0, but does not know if the value of $a_1$ is 1) ; $\neg K_1 l_0$ and $K_1 \neg l_1$ ($a_1$ does not know if the value of $a_0$ is 1 but it knows its value is 0)*

- *In $\gamma$ : $\neg K_0 \neg l_1$ and $K_0 l_0$ ($a_0$ knows its value, which is 0, but does not know if $a_1$ has value 0) ; $K_1 l_0$ and $K_1 l_1$ ($a_1$ knows both its value, which is 1, and the value of $a_0$ which is 1)*

*Knowledge is sparse though : one iteration of the knowledge semantics shows that in states $\alpha$ and $\beta$, $\neg K_1(K_0 \neg l_1)$, so that $a_1$ does not know if $a_0$ knows the value of $a_1$ in these two states (which is 1), and in $\beta$ and $\gamma$, $\neg K_0(K_1 l_0)$, that is, $a_0$ does not know if $a_1$ knows the value of $a_0$ in these two states (which is 0). One could guess that even starting with this restricted set of states, no agent can decide, coherently, to reach consensus. We will come back to this in more generality in Section 5.4.*

**Definition 8** (Language of action model logic)**.** $\varphi \in \mathcal{L}^{\text{stat}}_{KC\otimes}(A, P)$ and $\alpha \in \mathcal{L}^{\text{act}}_{KC\otimes}(A, P)$ are defined by:

$$\varphi ::= p \mid (\neg\varphi) \mid (\varphi \wedge \varphi) \mid K_a\varphi \mid [\alpha]\varphi$$
$$\alpha ::= (\mathsf{M}, \mathsf{s})$$

where $p \in P$, $a \in A$.

**Definition 9** (Semantics of formulas and actions)**.** *Epistemic state $(M, s)$ with $M = \langle S, \sim, L\rangle$, action model $\mathsf{M} = \langle \mathsf{S}, \sim, \mathsf{pre}\rangle$, and $\varphi \in \mathcal{L}^{\text{stat}}_{KC\otimes}(A, P)$ and $\alpha \in \mathcal{L}^{\text{act}}_{KC\otimes}(A, P)$.*

| | | |
|---|---|---|
| $M, s \models p$ | *iff* | $p \in L(s)$ |
| $M, s \models \neg\varphi$ | *iff* | $M, s \not\models \varphi$ |
| $M, s \models \varphi \wedge \psi$ | *iff* | $M, s \models \varphi$ *and* $M, s \models \psi$ |
| $M, s \models K_a\varphi$ | *iff* | *for all* $s' \in S : s \sim_a s'$ *implies* $M, s' \models \varphi$ |
| $M, s \models [\alpha]\varphi$ | *iff* | *for all* $M, \mathsf{M}, s, \mathsf{s} : M, s \models \mathsf{pre}(\mathsf{s})$ *implies* $(M \otimes \mathsf{M}, (s, \mathsf{s})) \models \varphi$ |

*The* restricted modal product $M \otimes \mathsf{M} = \langle S', \sim', L'\rangle$ *is defined as:*

| | | |
|---|---|---|
| $S'$ | $=$ | $\{(s, \mathsf{s}) \mid s \in S, \mathsf{s} \in \mathsf{S},$ *and* $M, s \models \mathsf{pre}(\mathsf{s})\}$ |
| $(s, \mathsf{s}) \sim'_a (t, \mathsf{t})$ | *iff* | $s \sim_a t$ *and* $\mathsf{s} \sim_a \mathsf{t}$ |
| $p \in L'(s, \mathsf{s})$ | *iff* | $p \in L(s)$ |

Now, we can organize Kripke models as a category, by defining Kripke model morphisms :

**Definition 10** (Morphism of Kripke models)**.** *Let $M = \langle S, \sim^A, L^{AP}\rangle$ and $N = \langle T, \sim^A, L^{AP}\rangle$ be two Kripke models. A morphism of Kripke models*

*is a morphism $f$ of the underlying Kripke frames of $M$ and $N$ such that $L^{AP}(f(s)) \subseteq L^{AP}(s)$ for all states $s$ in $S$.*

We note $\mathcal{KM}$ for the category of Kripke models with such morphisms. We note that $\mathcal{KM}$ is cartesian as is $\mathcal{K}$ the category of Kripke frames, with the cartesian product being defined as the Kripke model with underlying Kripke frame being the cartesian product of the underlying Kripke frames and the function $L^{AP}$ is defined on the product states $(s,t)$ by $L^{AP}(s,t) = L^{AP}(s) \cup L^{AP}(t)$.

We prove now that morphisms can only "lose knowledge" (whereas it is well-known that $p$-morphisms, [7] keep knowledge invariant) :

**Theorem 2.** *Consider now two Kripke models $M' = \langle S', \sim'^A, L' \rangle$ and $M = \langle S, \sim^A, L \rangle$, and a morphism $f$ from $M$ to $M'$. Then, $f$ "can only lose knowledge" for every agent $a$, i.e. for all $a \in A$, for all states $s \in M$, $M', f(s) \models K_a\phi \Rightarrow M, s \models K_a\phi$.*

*Proof.* Recall that :

$M, s \models K_a\varphi$ *iff* *for all* $s' \in S : s \sim_a s'$ *implies* $M, s' \models \varphi$ .

Consider $t \sim_a s$ : $f(t) \sim_a f(s)$ and as $M', f(s) \models K_a\phi$, by definition of the semantics of $K_a$ that we recapped above, we know that $M', f(t) \models \phi$. Therefore $\phi \in L^{AP}(f(t))$, and by definition of morphisms of Kripke models, $\phi \in L^{AP}(f(t)) \subseteq L^{AP}(t)$. So $M, t \models \phi$ and $M, s \models K_a\phi$. □

Note that knowledge is not defined in the topological approach to distributed computing [18]. We will establish below mappings between simplicial complexes and epistemic models, enabling us to prove topological assertions about epistemic models, and knowledge assertions about models based on complexes. The facets of simplicial complexes correspond to the states of epistemic models.

### 4.2 Action models for tasks

We now turn our attention to information change. Recall that an *action model* is a structure $\mathsf{M} = \langle \mathsf{S}, \sim, \mathsf{pre} \rangle$, where $\mathsf{S}$ is a domain of *action points*, such that for each $a \in A$, $\sim_a$ is an equivalence relation on $\mathsf{S}$, and $\mathsf{pre} : \mathsf{S} \to \mathcal{L}$ is a preconditions function that assigns a *precondition* $\mathsf{pre}(\mathsf{s})$ to each $\mathsf{s} \in \mathsf{S}$.

Each action can be thought of as an announcement made by the environment, which is not necessarily public, in the sense that not all system agents receive these announcements.

Let $M = \langle S, \sim^A, L^{AP} \rangle$ be a Kripke model and $A = \langle T, \sim, \mathsf{pre} \rangle$ be an action model. The *product update model* $M[A] = \langle S \times T, \sim^A, L^{AP} \rangle$, where

each world of $M[A]$ is a pair $(s,t)$ where $s \in S, t \in T$, such that $\mathsf{pre}(t)$ holds in $s$. Then, $(s,t) \sim_a (s',t')$ whenever it holds that $s \sim_a s'$ and $t \sim_a t'$. The valuation of $p$ at a pair $(s,t)$ is just as it was at $s$.

In this definition, the worlds of $M[A]$ are obtained by making copies of each world $s$ of $M$, one copy per event $t \in A$. A pair $(s,t)$ is to be included in the worlds of $M[A]$ if and only if $M$ satisfies the precondition $preA(t)$ of event $t$.[4].

We extend the task formalism in terms of Kripke frames from Section 3.2 to define an epistemic model. Recall that a task is a triple $\mathcal{T} = \langle \mathcal{G}_0, \mathcal{G}_f, \Delta \rangle$, where the set of initial initial states $\mathcal{G}_0$, is such that in every initial state, the environment is in the same state, and $\mathcal{G}_f$ is a set of states, with the same environment state. Also, $\Delta$ is a carrier morphism of Kripke frame $I = \langle \mathcal{G}_0, \sim^A \rangle$ to Kripke frame $O = \langle \mathcal{G}_f, \sim^A \rangle$. For the semantics, we assume that the local state of a process in either an input or an output state, is determined by the value of a variable. And just as an example simplify notation, we can assume these local variables take values from the set $\{0,1\}$. Thus, it is sufficient to consider atomic propositions $l_i$ interpreted as "$l_i$ is true when agent $a_i$ has value 1, otherwise it has value 0", as in Example 2. Then, the *input model* is $\langle \mathcal{G}_0, \sim^A, L^{AP} \rangle$, where $\langle \mathcal{G}_0, \sim^A \rangle$ is the input Kripke frame (see Definition 4), and the function $L^{AP} : S \to 2^{\mathrm{Lit}(AP)}$ is such that $l_i$ is true in a state $s$ precisely when the input variable of $a_i$ is 1.

In this case the indistinguishability relation $\sim^A$ matches the meaning of the atomic propositions, in the following sense. An agent $a_i$ knows the value of its input variable in a state $u$, because the atomic proposition $l_i$ has the same value in every $v$ such that $u \sim_i v$. More generally, one can infer $\sim^A$ from the values of the atomic propositions (and vice-versa, modulo renaming of names of atomic propositions).

The action model for the task, $\mathcal{T} = \langle \mathsf{S}, \sim, \mathsf{pre} \rangle$, is defined as follows. The action points, $\mathsf{S}$ is identified with $\mathcal{G}_f$, so action point $ac = \langle d_0, \ldots, d_n \rangle$ is interpreted as "agent $a_i$ decides value $d_i$", with precondition that is true in every input state $u$ such that $ac \in \Delta(u)$.

### 4.3 Action models in distributed computing

We assume that each state $s \in \mathcal{G}_0$ has the same environment state, $\epsilon$, representing that the shared memory is empty. To study tasks, assume the state of a process in $s$ is determined by the contents of a read-only local variable, which can take values from some finite set of values. Also, we

[4]Usually pointed Kripke models and action models are used, but we do not need them here.

assume the protocol $D$ is scheduler oblivious, to make the framework more uniform, see below.

Fix a model of distributed computation $M = (\mathcal{G}_0, P_e, \tau, \Psi)$, with initial states $\mathcal{G}_0 \subseteq \mathcal{G}$, and a deterministic protocol $D$, with some integer $N$, and consider the set $\mathcal{S}_N(D, M, \mathcal{G}_0)$ of $N$-admissible executions, where each process executes the same number of actions can be seen as a product of $\mathcal{G}_0$ and a set of composed scheduling actions $R$ as in Example 2.2.4, where each composed block action is of the form $sc_1 \odot sc_2 \odot \cdots \odot sc_R$.

The *action model for* $M$ is $A = \langle R, \sim, \mathsf{pre} \rangle$, where each action in $R$ is a composed block action of the form $sc_1 \odot sc_2 \odot \cdots \odot sc_R$, and $\mathsf{pre}$ is empty. That is, we allow each action to be applied to any initial state[5]. Then, $\sim$ is obtained directly from protocol frame in Section 3.3. Recall that the protocol Kripke frame can be written as $P = \langle S, \sim^A \rangle$, where $S = \mathcal{G}_0 \times AC$ the cartesian product defined in Lemma 1 between frame $\mathcal{G}_0$ and the action frame $AC$. More precisely, for two states $sc, sc' \in R$, it holds $sc \sim sc'$ whenever for two states $u, v \in \mathcal{G}_0$ with $u \sim_a v$, it holds that $u \odot sc \sim_a v \odot sc'$ in $P$.

For the following lemma, we assume the protocol $D$ is *scheduler oblivious,* in the sense that the protocol's behavior on two initial states with the same schedule is the same, except for possible differences in output values (as affected by input values). Equivalently, we could take the following lemma as an assumption. For distributed task computability purposes, the protocol $D$ can be assumed to be full-information (remember all the past, and send the whole local state in each write operation), which is certainly scheduler oblivious.

**Lemma 3.** *The relation $\sim$ of an action model is well defined. Namely, consider two $sc, sc' \in R$. If for two states $u, v \in \mathcal{G}_0$ with $u \sim_a v$, it holds that $u \odot sc \sim_a v \odot sc'$ in $P$, then for any two other states $u', v' \in \mathcal{G}_0$ with $u' \sim_a v'$ it also holds that $u' \odot sc \sim_a v' \odot sc'$.*

An action model $A = \langle R, \sim, \mathsf{pre} \rangle$ for a model $M$ with protocol $D$ and $N$ does indeed depend $M$, $D$ and $N$, although the action points $R$ as we defined may not seem to do so. An action point $ac$ in $R$ is typically just a sequence of sets of agents to be scheduled, and hence the same set $R$ can be used for different models. Thus, remarkably, the effect of different models is captured solely by $\sim$.

[5] A more precise semantics, could uses a copy of each block action for each input state, and specify a precondition that allows the copy to be applied only to that state; see discussion below on scheduler oblivious protocols.

We have defined an input epistemic model $I = \langle S, \sim^A, L^{AP} \rangle$ and an action model $A = \langle T, \sim, \mathsf{pre} \rangle$, and we get the product update model $I[A] = \langle S \times T, \sim^A, L^{AP} \rangle$, whose Kripke frame is exactly the protocol Kripke frame of Section 1, but decorated with the atomic propositions $l_i$. Moreover, a state in the product is of the form $(u, ac)$, where $u$ is an initial state in $\mathcal{G}_0$ and $ac$ is an action in $R$, and it has the same propositions that the input state $u$.

The action model for tasks is defined in an analogous way, and there is an epistemic model for the output Kripke frame by the corresponding epistemic product.

In the important following result, we lift the task solvability condition of Theorem 1 from the category of Kripke frames to the category of Kripke models :

**Theorem 3.** *Suppose we have input, output and specification Kripke models, respectively $I$, $O$ and $\Delta \subseteq I \times O$. Then task solvability in the sense of Definition 5 is equivalent to the existence of a morphism $h$ such that the same equation as in Theorem 1 holds :*

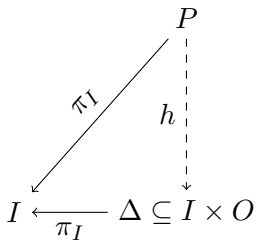


By Theorem 2, we know that the knowledge of each agent (or process) can only decrease (or stay constant) along the two (surjective) arrows $\pi_I$ of this diagram. So the diagram above has a simple an illuminating interpretation : we can only improve knowledge from $I$ to $P$ (the protocol should improve knowledge of all agents through communication). Now the task is solvable if and only if there is enough knowledge from all agents (map $h$) to that to have at least the knowledge specified by Kripke model $\Delta$.

# 5 Combinatorial topology and Kripke models

We now describe the equivalence of categories between (proper) Kripke frames and (pure chromatic) simplicial complexes; the semantics of distributed systems is then transported from Kripke models to simplicial complexes

(suitably decorated). This shows that "knowledge" exhibits topological invariants.

## 5.1 Simplicial complexes and simplicial models

We consider now an apparently very different structure from a Kripke graph, which has also been used to model distributed systems [18] (for a textbook covering combinatorial topology notions see [26]) :

**Definition 11** (Simplicial complex)**.** *A simplicial complex $C$ is a family of non-empty finite subsets of a set $S$ such that for all $X \in C$, $Y \subseteq X$ implies $Y \in C$ ($C$ is downwards closed).*

Elements of $S$ (identified with singletons) are called vertices, elements of $C$ of greater cardinality are called faces. The dimension of a face $X \in C$, $\dim X$, is the cardinality of $X$ minus one. The maximal faces of $C$ (i.e. faces that are not subsets of any other face) are called facets. The dimension of a simplicial complex is the maximal dimension of its faces. Pure simplicial complexes are simplicial complexes such that the maximal faces are all of the same dimension.

**Definition 12** (Simplicial maps)**.** *Let $C$ and $D$ be two simplicial complexes. A simplicial map $f : C \to D$ is a function that maps the vertices of $C$ to the vertices of $D$ such that for all faces $X$ of $C$, $f(C)$ (the image set on the subset of vertices $C$) is a face of $D$.*

Now, we can define pure simplicial maps respecting facets. Finally, we will associate colors to each vertex of simplicial complexes, representing, as in [18], the names of the different processes involved in a protocol. We also define chromatic simplicial maps as the simplicial maps respecting colors.

**Definition 13** (Pure simplicial complexes)**.** *A pure simplicial complex $C$ is a simplicial complex such that all facets have the same dimension as $C$.*

**Definition 14** (Pure simplicial maps)**.** *Let $C$ and $D$ be two pure simplicial complexes. A pure simplicial map $f : C \to D$ is a simplicial map such that $\dim f(X) = \dim X$, for all $X \in C$, i.e. is such that $f$, restricted to vertices of a given simplex, is injective.*

**Definition 15** (Chromatic simplicial complex)**.** *A chromatic simplicial complex $C$ on a vertex set $S$, colored by a set of colors $L$, is a simplicial complex together with a (coloring) map $l : S \to L$ (inducing a map, still called $l$, from $2^S$ to $2^L$) such that for all $X \in C$, $\dim l(X) = \dim X$.*

The condition of the coloring map is indeed to be interpreted as the property that the vertices of $X$ have all distinct colors.

**Definition 16** (Chromatic simplicial maps)**.** *A chromatic simplicial map is a simplicial map $f : C \to C'$ which preserves the coloring, i.e. such that $l'(f(X)) = l(X)$, for all $X \in C$.*

Note that a chromatic simplicial map is necessarily injective on each simplex, since as it preserves colors, and points of any simplex are of distinct colors, $f$ cannot equate two points of the same simplex. This also implies that for all $X$ and $Y$ chromatic simplicial complexes, $f(X \cap Y) = f(X) \cap f(Y)$.

Let $p\mathcal{CS}$ be the category of pure chromatic simplicial complexes.

A common way to describe actual states in combinatorial topology is by decorating vertices of simplicial complexes by local states of processes. Instead of describing actual values, we rely on a logical languages to describe the properties of local states.

Let $AP$ be a countable set of propositional variables and $A$ a set of agents. We use the same notations as in Section 4 : the set of *literals* over $AP$ is $\mathrm{Lit}(AP) = AP \cup \{\neg p \mid p \in AP\}$. If $X \subseteq \mathrm{Lit}(AP)$, then $\overline{X} = \{\overline{\ell} \mid \ell \in X\}$; $X$ is *consistent* iff $\forall\ \ell \in X$, $\overline{\ell} \notin X$; and $X$ is *$AP$-maximal* iff $\forall\ p \in AP$, either $p \in X$ or $\neg p \in X$.

**Definition 17.** *A simplicial model is $(C, l, v)$ where $(C, l)$ is a pure chromatic simplicial set, and $v : S \to \wp(AP)$ an assignment of subsets of the set of literals of $AP$ for each state $s \in S$ such that for all facets $f = (s_0, \ldots, s_n) \in C$, $\bigcup_{i=0}^{n} v(s_i)$ (that we denote as $v(f)$ by an abuse of notation) is $AP$-maximal.*

## 5.2 Equivalence between (proper) Kripke models and simplicial models

We have the main result of this section :

**Theorem 4.** *Let $A$ be a finite set and $p\mathcal{CS}_A$ (resp. $\mathcal{K}_A$) be the full subcategory of pure chromatic simplicial complexes with colors in $A$ (resp. the full subcategory of proper Kripke frames with agent set $A$). $p\mathcal{CS}_A$ and $\mathcal{K}_A$ are equivalent categories.*

*Proof.* We construct functors $F : p\mathcal{CS} \to \mathcal{K}$ and $G : \mathcal{K} \to p\mathcal{CS}$ as follows.

Let $C$ be a pure chromatic complex on the set of colors $A$. We associate $F(C) = \langle S, \sim^A \rangle$ with $S$ being the set of facets of $C$ and the equivalence relation $\sim_a$, for all $a \in A$, generated by the relations $X \sim_a Y$ (for $X$ and $Y$ facets of $C$) if $a \in l(X \cap Y)$.

For $f : C \to D$ a morphism in $p\mathcal{CS}$, consider facets $X$ and $Y$ of $C$ such that $X \sim_a Y$ in $F(C)$, thus $a \in l(X \cap Y)$. Because $f$ is a chromatic simplicial map, we have seen that $f(X \cap Y) = f(X) \cap f(Y)$. Because $f$ is chromatic, $f(X)$ is a facet of $D$ which has the same colors as $X$, hence $a \in l(f(X))$. Similarly for $f(Y)$, which is such that $a \in l(f(Y))$. The colors of $f(X)$ (resp. $f(Y)$) are in bijection with the vertices of $f(X)$ (resp. $f(Y)$), hence $a \in l(f(X) \cap f(Y)$, therefore $f(X) \sim_a f(Y)$.

Conversely, consider a Kripke frame $M = \langle S, \sim^A \rangle$. Define $G(M)$ to be the quotient of the coproduct of $n$-simplexes $\coprod_{s \in S} \{v_0^s, \ldots, v_n^s\}$ where $l(v_i^s) = a_i \in A$ ($A = \{a_0, \ldots, a_n\}$) by the relation $R$ defined as : $v_i^s R v_i^{s'}$ if and only if $s \sim_i s'$. This relation is then extended on higher simplices by $\{x_0, \ldots, x_k\} R \{y_0, \ldots, y_l\}$ if and only if all $x_i$ are in relation (by $R$) with some $y_{j_i}$ and inversely, all $y_j$ are in relation with some $x_{i_j}$. $R$ can thus be seen as a sub-chromatic simplicial complex of $X \times X$ and the quotient is well-defined, as a chromatic simplicial complex. It is pure since, as we equate only vertices with the same colors, and that colors are in bijection with extremal points in all simplexes, we cannot equate a simplex with a lower dimensional simplex. The facets are the $\{v_0^s, \ldots, v_n^s\}$ where $s \in S$, since the Kripke frames we consider being proper, we cannot equate two facets together.

Consider now a Kripke frame $M = \langle S, \sim^A \rangle$ in $\mathcal{K}$ with agent set $A$. $FG(M)$ is the Kripke frame $N = \langle T, \sim'^A \rangle$ such that $T$ is the set of facets of $G(M)$. But we have just seen that the set of facets of $G(M)$ are the sets of vertices $\{v_0^s, \ldots, v_n^s\}$ (where $s \in S$), therefore, is in bijection with $S$. Finally, in $FG(M)$, $p \sim'_a q$ if and only if $a \in l(p \cap q)$, where $l$ is the coloring, in $G(M)$, of $p$ and $q$ which are facets in $G(M)$. But facets in $G(M)$ are just in direct bijection with the worlds of $M$, i.e. $p = \{v_0^s, \ldots, v_n^s\}$ and $q = \{v_0^t, \ldots, v_n^t\}$ where $s, t \in M$. Note that $l(v_i^s) = a_i$ and $l(v_i^t) = a_i$ so $a \in l(p \cap q)$ means that some $a = a_i$ for some $i$ and $v_i^s R v_i^t$. This can only be the case, by definition of $G(M)$ if $s \sim_i t$. This proves that $FG(M)$ and $M$ as isomorphic Kripke frame.

Consider now a pure chromatic simplicial complex $X \in p\mathcal{CS}$. It is easily seen that $GF(X)$ is isomorphic, as a pure chromatic simplicial complex, to $X$, hence $p\mathcal{CS}$ and $\mathcal{K}$ are equivalent categories. □

We write $\mathbf{K}$ (resp. $\mathbf{f} : \mathbf{K} \to \mathbf{L}$) when $K$ is a Kripke frame for the corresponding pure chromatic simplicial complex (resp. when $f$ is a morphism of Kripke frame, for the corresponding simplicial map).

Let $A$ be a finite set (of "agents") and $\mathcal{SM}_{A,AP}$ (resp. $\mathcal{KM}_{A,AP}$) be the

full subcategory of simplicial models with colors in $A$ and vertices decorated with formulas in $AP$ (resp. the full subcategory of proper Kripke models with agent set $A$ and atomic propositions in $AP$). Theorem 4 extends to the following theorem, in a straightforward manner :

**Theorem 5.** *$\mathcal{SM}_{A,\mathcal{G}}$ and $\mathcal{KM}_{A,\mathcal{G}}$ are equivalent categories.*

In the category of pure chromatic simplicial complexes (resp. simplicial models), the cartesian product of $K = (C, S, l)$ with $K' = (C', S', l')$ (resp. simplicial models $K = (C, S, l, v)$ and $K' = (C', S', l', v')$), where $C$ is a family of non-empty finite subsets of $S$, $C'$ a family of non-empty subsets of $S'$, $l$ is a labelling map from $S$ to $L$, and $l'$ is a labelling map from $S'$ to $L$, is defined as $K \times K' = (D, S \times S', m)$ (resp. $K \times K' = (D, S \times S', m, w)$),

- $D = \{\{(s_0, s'_0), \ldots, (s_k, s'_k)\} | \{s_0, \ldots, s_k\} \in C, \{s_0, \ldots, s_k\} \in C', l(s_i) = l'(s'_i)\}$
- $m(s_i, s'_i) = l(s_i) = l'(s'_i)$ (resp. $w$ is defined by $w(s) = v(s) \cup v'(s)$)

Of course, there is a simplicial map $\pi_K$ from $K \times K'$ to $K$ and a map $\pi_{K'}$ from $K \times K'$ to $K'$ (which are the first and second projection on each simplex, mapped onto subsets of simplices). And of course, whenever we have a morphism $u$ in $p\mathcal{CS}$ (resp. $\mathcal{SM}$) from $U$ to $K$ and another one $v$ from $U$ to $K'$, the map $h$ from $U$ to $K \times K'$ defined by $h(x) = (u(x), v(x))$ is such that $u = \pi_K \circ h$ and $v = \pi_{K'} \circ h$. These are the direct translations under the equivalences of Theorem 4 (resp. 5) of the cartesian product of Kripke frames (resp. Kripke models).

**Example 3.** *We have pictured below the product $\mathbf{I} \times \mathbf{O}$, with its two projections $\pi_{\mathbf{O}}$ and $\pi_{\mathbf{I}}$, for the input complex $\mathbf{I}$ for two processes, having boolean input values, and for the output complex $\mathbf{O}$ for binary consensus :*

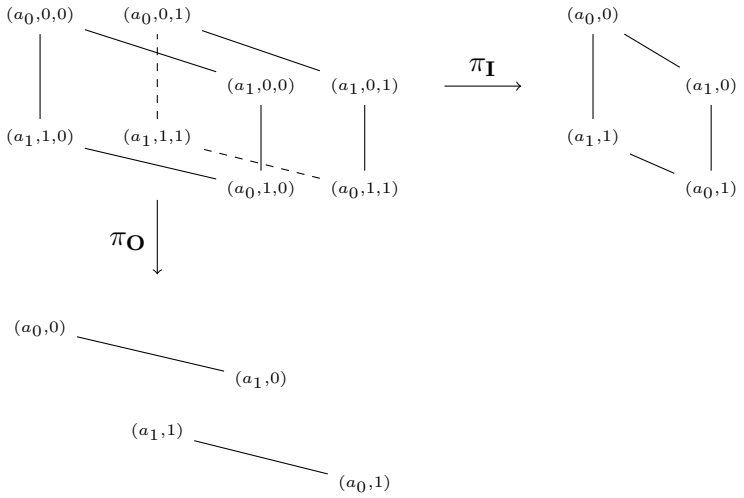


*In the example above the vertices of the simplicial models are decorated with predicates, as follows : for $O$, $(a_i, j)$ is a notation for, agent $a_i$ is such that "its output value is $j$" (this is the predicate that we will call $d_j$ in the sequel) ; for $I$, $(a_i, j)$ is a notation for, agent $a_i$ is such that "its input value is $j$" (this is the predicate that we will call $l_j$ is the sequel). Thus in the product, $(a_i, j, k)$ means, agent $a_i$ is such that "its input value is $j$" and "its output value is $k$" (as expected for a cartesian product, this is the set of predicates $\{l_j, d_k\}$).*

In fact, there is much more categorical structure in pure chromatic simplicial complexes, although to see this, we would have to go into much more involved arguments, that we do not need in this simplified presentation. Let us just mention that the overall categorical structure can be derived by observing that simplicial complexes form a quasi-topos with all (small) limits and (small) colimits [4], and that $p\mathcal{CS}$ can be seen as a particular subcategory of a slice category of these simplicial complexes.

### 5.3 Action models, protocol complexes and task solvability

Given an action model $\mathsf{M} = \langle \mathsf{S}, \sim, \mathsf{pre} \rangle$, with equivalent simplicial model $\mathsf{M}$, and for any input Kripke model $I$, with equivalent simplicial model $\mathbf{I}$, we form the functor $P$ which associates the sub-Kripke frame of the cartesian product $I \times \mathsf{M}$ (where some of the cartesian product has been filtered out due to preconditions), as in Section 3.3.

**Example 4.** *In the atomic read/write model for two processes $a_0$ and $a_1$ and 1-admissible runs,* $\mathbf{M}$ *corresponds to the three possible schedules of executions, with no precondition i.e. $P$ is the cartesian product with $\mathsf{M}$ and is the unique endofunctor on pure chromatic simplicial complexes, commuting with colimits, such that its image on a segment* $^{a_0 \ — \ a_1}$ *is the chromatic subdivision of the segment* $^{a_0 \ - \ a_1 \ - \ a_0 \ - \ a_1}$ *. In the Figure below, we show the effect of $P$ on the input complex* $\mathbf{I}$ *of Example 5, together with the structure map $\pi_I$ from $P(\mathbf{I})$ to* $\mathbf{I}$ *:*

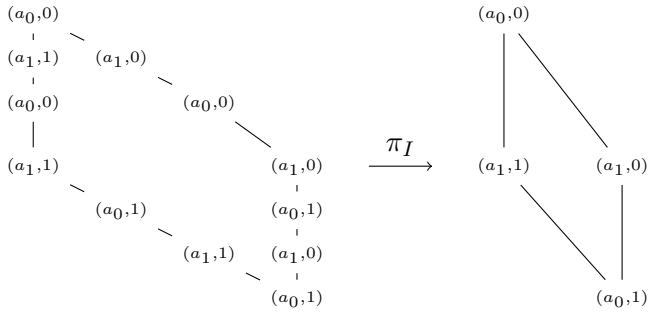


By the equivalence of categories of Theorem 5, this creates a functor that we still write $P$ associating some subcomplex of the cartesian product $\mathbf{I} \times \mathsf{M}$. Of course, since this is a sub-object of a cartesian product, we have a canonical structure map $\pi_I : P(\mathbf{I}) \to \mathbf{I}$.

We recall the classical approach in combinatorial topology for task solvability, see e.g. [18] :

**Definition 18.** *A decision task is a triple $T = (\mathcal{I}, \mathcal{O}, \Phi)$ where $\mathcal{I}, \mathcal{O}$ are $n$-chromatic complexes and $\Phi$ is a map which satisfies:*

1. *$\Phi(\sigma)$ is a subcomplex of $\mathcal{O}$.*
2. *$\Phi(\tau) \subseteq \Phi(\sigma)$ if $\tau < \sigma$.*
3. *$id(\sigma) = id(\Phi(\sigma))$.*

In general, a map $\Phi$ from a simplicial complex $\mathcal{C}$ to a complex $\mathcal{D}$ which satisfies the conditions *1* and *2* of Definition 18 is called a *carrier map*. In Definition 18 each vertex represents the state of a single process. A simplex $\sigma^{(k)}$ is used to represent compatible states, of $k+1$ processes. In addition $\Phi(\sigma)$ defines which output states are legal for each input simplex $\sigma \in \mathcal{I}$. For

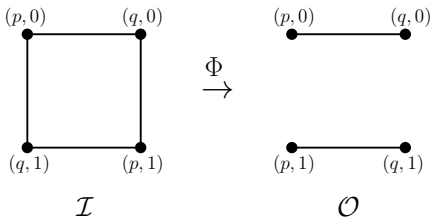


Figure 1: The binary consensus task.

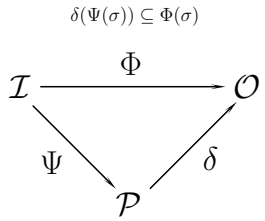


Figure 2: Solvability of a task $(\mathcal{I}, \mathcal{O}, \Phi)$ by a protocol $\mathcal{P}$ with decision map $\delta$

instance in the binary consensus task each process proposes a binary value and they have to agree in one of them. For two processes $p$ and $q$ there are four possible input configurations, which are represented as the maximal edges (maximal simplices) of $\mathcal{I}$ in Figure 1. In this case if $\sigma_1 = \{(p,0),(q,0)\}$ then $\Phi(\sigma_1) = \sigma_1$ or $\Phi(\sigma_2) = \mathcal{O}$ if $\sigma_2 = \{(p,0),(p,1)\}$.

A chromatic complex $\mathcal{P}$ *solves a task* $T = (\mathcal{I}, \mathcal{O}, \Phi)$ if there exists a simplicial map $\delta$ from $\mathcal{P}$ to $\mathcal{O}$ such that

1. $\delta(\Psi(\sigma)) \subseteq \Phi(\sigma)$ for all simplex $\sigma$.
2. $id(v) = id(\delta(v))$ for all vertex $v \in \mathcal{P}$.

where $\Psi$ is a carrier map from $\mathcal{I}$ to $\mathcal{P}$, see Figure 2. The Asynchronous Computability Theorem states that a decision task $T = (\mathcal{I}, \mathcal{O}, \Phi)$ is solvable by a wait-free protocol using read/write memory if only if there exist a chromatic subdivision $\chi^{(k)}(\mathcal{I})$ such that solves $T$.

Now, it is easy to see that, by Theorem 4, this is exactly equivalent to the formulation, using Kripke frames, of Section 3.2.

As a matter of fact, quite generally speaking, a carrier map $\Phi$ can be seen as a functor from $p\mathcal{CS}$ to $p\mathcal{CS}$ that is such that $P \circ \Phi = Id_{p\mathcal{CS}}$. Indeed, restricting the latter equality on simplexes $\sigma$ of some input complex $\mathbf{I}$ implies

that $\Phi(\sigma)$ is the subcomplex $P(\mathbf{I})$. We can interpret the task solvability diagram of Section 3.3 in purely topological terms, this interest being that some topological invariants will prevent us from finding a map $h$ as above, showing impossibility of the corresponding task specification. As $\pi_I$, $h$ and $\pi_I$ are simplicial maps, we should have a corresponding diagram using any topological invariant functor, such as homology or homotopy functors (see e.g. Example 5 below). In case of atomic read/write memory models, we know that $P(I)$ corresponds to some subdivision of $I$, hence $\pi_I : P(I) \to I$ is a weak homotopy equivalence (which in turn implies knowledge is not improving in $P(I)$ with respect to $I$), this restricts a lot what task specifications $\Delta$ can be solved, as also exemplified in Example 5.

**Example 5.** *In the case of binary consensus, the task specifications is given by the Kripke frame $\Delta$ below. The input Kripke model is composed of the four states which account for all binary inputs on two processes (modelled by the two predicates, $l_0$ for describing the input of agent $a_0$ and $l_1$ for describing the input of agent $a_1$). The output complex is composed of the two states where the two local values agree (we renamed the corresponding predicates for $a_i$, $d_i$) :*

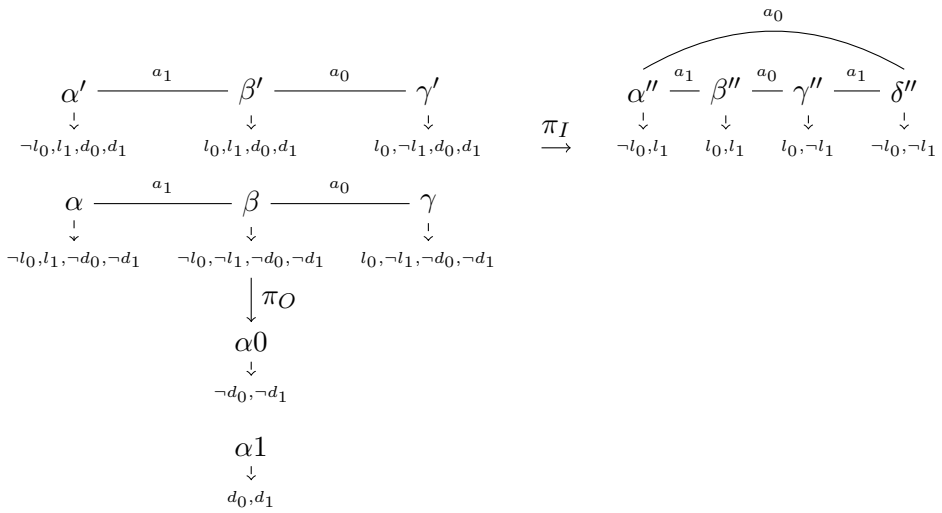


*Using the equivalence of categories with simplicial models, we get $\mathbf{\Delta} \subseteq \mathbf{I} \times \mathbf{O}$, where the state $((a_0, 1, 0), (a_1, 1, 0))$ (and $((a_0, 0, 1), (a_1, 0, 1))$) corresponding to both processes beginning with input 1 and deciding 0 (and corresponding to both processes beginning with input 0 and deciding 1), have*

*been left out, shown below with its two projections on the input and output complexes :*

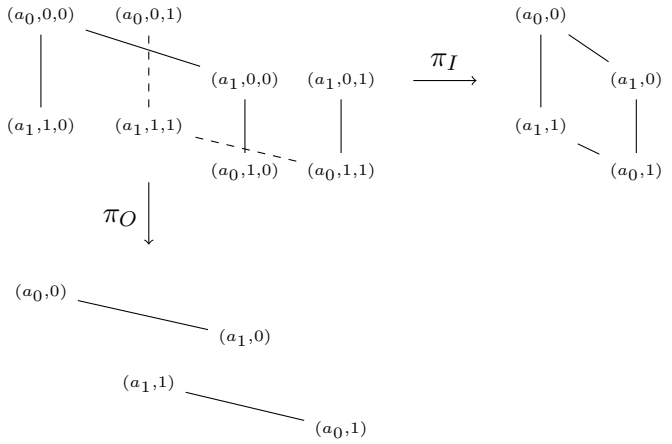


*In terms of topological invariants we see in the diagram above that the first homology group of $P(\mathbf{I})$ and of $\mathbf{I}$ are $\mathbb{Z}$ (they are both homotopically equivalent to the circle) and that $\pi_I$ induces the identity map in homology, whereas the first homology group of $\mathbf{\Delta}$ is 0. There is no factorization of the identity through the 0 map (which is $\pi_{\mathbf{I}}$ in homology), so binary consensus is not solvable.*

*Now the pseudo consensus specification is as follows :*

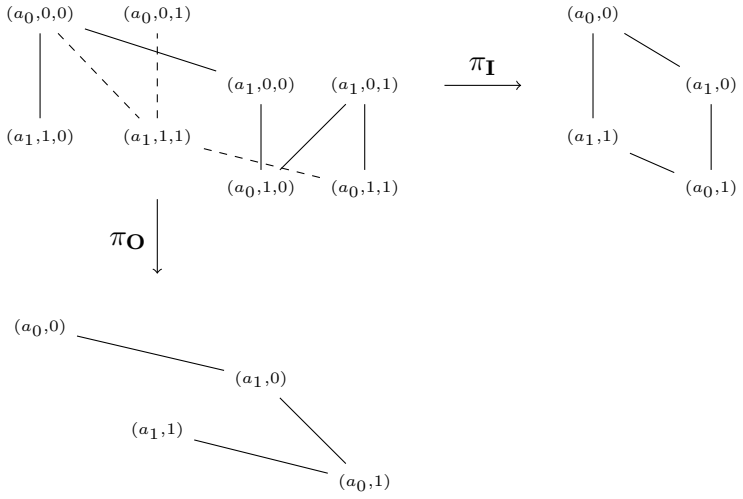

*Topologically, $P(\mathbf{I})$, $\mathbf{\Delta}$ and $\mathbf{I}$ are homotopy equivalent to a circle, and $\pi_I$ (the identity map in homology) can be factored through the homology of $P(\mathbf{I})$. Indeed, binary pseudo-consensus is solvable in one round in atomic read/write memory.*

### 5.4 Common knowledge and connectivity

For $B$ a subgroup of agents (or processes), we recall, from Section 4 the notion of group knowledge :

$$E_B\varphi = \bigwedge_{b\in B} K_b\varphi$$

Common knowledge for group $B$ is, semantically, the least solution to the equation :

$$C_B\varphi = \varphi \wedge E_B(C_B\varphi)$$

Given a Kripke model $M = \langle S, \sim^A, L^{AP}\rangle$, and a group of agents $B = \{s_1, \ldots, s_k\}$. Define $\sim_B$ to be the transitive closure of all the $\sim_a$, $a \in B$. This means that $s \sim_B t$ if and only if there exists $s_1, \ldots, s_l \in S$ such that $s \sim_{a_{i_1}} s_1 \sim_{a_{i_2}} s_2 \ldots s_l \sim_{a_{i_{l+1}}} t$ with $1 \leq a_{i_1}, \ldots, a_{i_{l+1}} \leq k$.

Now we have the following semantic characterization of $C_B\varphi$ at a state $s$ of a Kripke model $M = \langle S, \sim^A, L^{AP}\rangle$ :

$M, s \models C_B\varphi$ if and only if $M, t \models \varphi$ for all $t$ such that $s \sim_B t$.

In the sequel, by an abuse of notation, we will be identifying any Kripke model $M$ with its simplicial model counterpart $(S, l, v)$ under the equivalence of Theorem 5 between Kripke models and simplicial models, by identifying states $S$ of Kripke models with facets of the simplicial model $(S, l, v)$. Given $B$ a group of agents, we form the Kripke model (and equivalently the simplicial model) $M_B$ from $M$ by restricting the accessibility relation to those $\sim_a$ for $a \in B$, and the formulas associated to each state to contain only formulas describing agent in $B$. Then we have :

**Theorem 6.** *The formula $\varphi$ is common knowledge at state $s$ in model $M$ if and only if $\varphi$ is true on all facets in the same connected components of simplicial model $M$ restricted to the subgroup $B$ of agents. In other words : for all facets $t$ in $M_B$ which are in the same connected component as $s$ in $M_B$, we have $\varphi \in v(t)$.*

**Example 6.** *Let us consider again the binary consensus as specified in Example 5. In $\Delta$, there are two connected components : in the first one, by Theorem 6, it is common knowledge that $\neg l_0 \vee \neg l_1$, meaning that either $a_0$*

*started with* 0 *or* $a_1$ *started with 0. In the second component, by the same theorem, it is common knowledge that* $l_0 \vee l_1$*, meaning that either* $a_0$ *started with 1, or* $a_1$ *started with 1. Now, in the protocol complex of Example 4, there is just no common knowledge involving local input values of* $a_0$ *nor* $a_1$*. By Theorem 3, if binary consensus were solvable, we would have a morphism of Kripke model h from the protocol Kripke model P to* $\Delta$*, implying by Theorem 2 that any knowledge available in* $\Delta$ *would be available in P, in particular common knowledge. This is clearly not true here and binary consensus is not solvable.*

# 6 Conclusions

We have made a first step into defining a version of muti-agent dynamic epistemic logic using as models higher dimensional simplical complexes. Although the step is modest, it already shows that under a class of action models, topological invariants of the initial epistemic model are fully preserved. We have worked out in detail the class of action models for a well studied distributed computing setting, where asynchronous processes that can crash communicate with each other through read/write shared variables. We established a categorical equivalence between such systems and dynamic epistemic models, and hence a precise, close relationship between the combinatorial topology theory of distributed computing and dynamic epistemic logic.

Many interesting questions are left for future work. We have developed all our theory on facets, interpreting only the top dimensional simplexes used in distributed computing, which is sometimes what is done, but often also simplexes of lower dimension are used to model process crashes. Another, main avenue that we left for future study is the role of bisimulations, which are very important in dynamic epistemic logic, and have also been considered in algebraic topology. Of course it would be of interest to study other distributed computing settings, especially those which have stronger communication objects available in modern multiprocessor architecture, and which are known to yield complexes that preserve less well the topology of the input complex; indeed, their additional power is expressed as higher dimensional "holes" in the protocol complex. It would be very interesting to find a formalization of such topological properties in terms of knowledge, and thus obtain a generalization from common knowledge (that is tightly related to 1-dimensional connectivity) to other form of group knowledge (related to higher-dimensional connectivity).

**Acknowledgements.** We thank Carlos Velarde and David Rosenblueth for their involvement in the early stages of this research, and their help in developing the dual of a Kripke graph. This work was partially supported by PAPIIT-UNAM IN109917. Sergio Rajsbaum would like to acknowledge the Ecole Polytechnique for financial support through the 2016-2017 Visiting Scholar Program.

# References

[1] Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic Snapshots of Shared Memory. *J. of the ACM.*, 40(4):873–890, September 1993.

[2] Hagit Attiya and Sergio Rajsbaum. The combinatorial structure of wait-free solvable tasks. *SIAM J. Comput.*, 31(4):1286–1313, April 2002.

[3] Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations, and Advanced Topics.* Wiley, 2 edition, 2004. URL: `http://www.wiley.com/WileyCDA/WileyTitle/productCd-0471453242.html`.

[4] J. C. Baez and A. E. Hoffnung. Convenient Categories of Smooth Spaces. *ArXiv e-prints*, July 2008. `arXiv:0807.1704`.

[5] A. Baltag, L.S. Moss, and S. Solecki. The logic of common knowledge, public announcements, and private suspiciouns. In *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)*, pages 43–56, 1998.

[6] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139(2):165–224, 2004. `doi:10.1023/B:SYNT.0000024912.56773.5e`.

[7] Alexandru Baltag and Bryan Renne. Dynamic epistemic logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy.* Metaphysics Research Lab, Stanford University, winter 2016 edition, 2016.

[8] N. Bezhanishvili and W. Van der Hoek. Structures for Epistemic Logics. In Outstanding Contributions to Logic, Volume 5, pages 175–202, 2014.

[9] Ofer Biran, Shlomo Moran, and Shmuel Zaks. A Combinatorial Characterization of the Distributed 1-Solvable Tasks. *J. Algorithms*, 11(3):420–440, 1990.

[10] Elizabeth Borowsky and Eli Gafni. Generalized flp impossibility result for t-resilient asynchronous computations. In *Proc. 25th Annual ACM Symp. on Theory of Computing*, STOC, pages 91–100, New York, NY, USA, 1993. ACM. URL: `http://doi.acm.org/10.1145/167088.167119`, `doi:10.1145/167088.167119`.

[11] Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computation (extended abstract). In *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '97, pages 189–198, New York, NY, USA, 1997. ACM.

[12] Elizabeth Borowsky, Eli Gafni, Nancy Lynch, and Sergio Rajsbaum. The BG distributed simulation algorithm. *Distributed Computing*, 14(3):127–146, 2001.

[13] Soma Chaudhuri. More choices allow more faults: set consensus problems in totally asynchronous systems. *Information and Computation*, 105(1):132–158, 1993.

[14] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer Publishing Company, Incorporated, 1st edition, 2007.

[15] M. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility Of Distributed Commit With One Faulty Process. *Journal of the ACM*, 32(2), April 1985.

[16] Eli Gafni, Petr Kuznetsov, and Ciprian Manolescu. A generalized asynchronous computability theorem. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, PODC '14, pages 222–231, New York, NY, USA, 2014. ACM. URL: `http://doi.acm.org/10.1145/2611462.2611477`, `doi:10.1145/2611462.2611477`.

[17] John Havlicek. Computable obstructions to wait-free computability. *Distributed Computing*, 13(2):59–83, 2000.

[18] Maurice Herlihy, Dmitry Kozlov, and Sergio Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Elsevier-Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 2013.

[19] Maurice Herlihy and Sergio Rajsbaum. Simulations and reductions for colorless tasks. In *Proceedings of the 2012 ACM symposium on Principles of distributed computing*, PODC '12, pages 253–260, New York, NY, USA, 2012. ACM.

[20] Maurice Herlihy and Sergio Rajsbaum. The topology of distributed adversaries. *Distributed Computing*, 26(3):173–192, 2013.

[21] Maurice Herlihy and Nir Shavit. The topological structure of asynchronous computability. *J. ACM*, 46(6):858–923, November 1999. URL: `http://doi.acm.org/10.1145/331524.331529`, `doi:10.1145/331524.331529`.

[22] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Elsevier, 1 edition, 2012. URL: `https://www.elsevier.com/books/the-art-of-multiprocessor-programming-revised-reprint/herlihy/978-0-12-397337-5`.

[23] Maurice P. Herlihy. Impossibility and universality results for wait-free synchronization. In *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing*, PODC '88, pages 276–290, New York, NY, USA, 1988. ACM. URL: `http://doi.acm.org/10.1145/62546.62593`, `doi:10.1145/62546.62593`.

[24] Gunnar Hoest and Nir Shavit. Towards a topological characterization of asynchronous complexity. In *Proc. 16th ACM Symp. Principles of distributed computing*, PODC, pages 199–208, New York, NY, USA, 1997. ACM.

[25] Alex Kogan and Erez Petrank. A methodology for creating fast wait-free data structures. In *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '12, pages 141–150, New York, NY, USA, 2012. ACM. URL: `http://doi.acm.org/10.1145/2145816.2145835`, `doi:10.1145/2145816.2145835`.

[26] Dmitry Kozlov. *Combinatorial Algebraic Topology*. Springer, 2007.

[27] M. C. Loui and H. H. Abu-Amara. *Memory requirements for agreement among unreliable asynchronous processes*, volume 4, pages 163–183. JAI press, 1987.

[28] Nancy Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996. URL: `https://www.elsevier.com/books/distributed-algorithms/lynch/978-1-55860-348-6`.

[29] Yoram Moses and Sergio Rajsbaum. A layered analysis of consensus. *SIAM J. Comput.*, 31(4):989–1021, April 2002. URL: http://dx.doi.org/10.1137/S0097539799364006, doi:10.1137/S0097539799364006.

[30] J.A. Plaza. Logics of public communications. In M.L. Emrich, M.S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.

[31] Timothy Porter. Interpreted systems and Kripke models for multiagent systems from a categorical perspective. In Theoretical Computer Science, Volume 323, Number 1, pages 235–266, 2004.

[32] Sergio Rajsbaum. Iterated Shared Memory Models. In *LATIN*, volume 6034 of *Lecture Notes in Computer Science*, pages 407–416. Springer, 2010.

[33] Sergio Rajsbaum, Michel Raynal, and Corentin Travers. The iterated restricted immediate snapshot model. In *COCOON*, volume 5092 of *Lecture Notes in Computer Science*, pages 487–497. Springer, 2008.

[34] Michel Raynal. Wait-free computing: An introductory lecture. *Future Gener. Comput. Syst.*, 21(5):655–663, May 2005. URL: http://dx.doi.org/10.1016/j.future.2004.05.005, doi:10.1016/j.future.2004.05.005.

[35] Yoram Moses Ronald Fagin, Joseph Halpern and Moshe Vardi. *Reasoning About Knowledge*. MIT Press, 1 edition, 1995. URL: https://mitpress.mit.edu/books/reasoning-about-knowledge.

[36] Michael Saks and Fotios Zaharoglou. Wait-Free k-Set Agreement is Impossible: The Topology of Public Knowledge. *SIAM J. Comput.*, 29(5):1449–1483, 2000. URL: http://dx.doi.org/10.1145/167088.167122, doi:10.1145/167088.167122.